

# Robust Shift-and-Invert Preconditioning: Faster and More Sample Efficient Algorithms for Eigenvector Computation \*

Chi Jin UC Berkeley chijin@eecs.berkeley.edu	Sham M. Kakade University of Washington sham@cs.washington.edu	Cameron Musco MIT cnmusco@mit.edu
Praneeth Netrapalli Microsoft Research, New England praneeth@microsoft.com	Aaron Sidford Microsoft Research, New England asid@microsoft.com	

November 2, 2015

## Abstract

In this paper we provide faster algorithms and improved sample complexities for approximating the top eigenvector of a matrix  $\mathbf{A}^\top \mathbf{A}$ . In particular we give the following results for computing an  $\epsilon$  approximate eigenvector - i.e. some  $x$  such that  $x^\top \mathbf{A}^\top \mathbf{A} x \geq (1 - \epsilon) \lambda_1(\mathbf{A}^\top \mathbf{A})$ :

- **Offline Eigenvector Estimation:** Given an explicit matrix  $\mathbf{A} \in \mathbb{R}^{n \times d}$ , we show how to compute an  $\epsilon$  approximate top eigenvector in time  $\tilde{O}\left(\left[\text{nnz}(\mathbf{A}) + \frac{d \text{sr}(\mathbf{A})}{\text{gap}^2}\right] \cdot \log 1/\epsilon\right)$  and  $\tilde{O}\left(\left[\frac{\text{nnz}(\mathbf{A})^{3/4} (d \text{sr}(\mathbf{A}))^{1/4}}{\sqrt{\text{gap}}}\right] \cdot \log 1/\epsilon\right)$ . Here  $\text{sr}(\mathbf{A})$  is the stable rank,  $\text{gap}$  is the multiplicative gap between the largest and second largest eigenvalues, and  $\tilde{O}(\cdot)$  hides log factors in  $d$  and  $\text{gap}$ . By separating the gap dependence from  $\text{nnz}(\mathbf{A})$  our first runtime improves classic iterative algorithms such as the power and Lanczos methods. It also improves on previous work separating the  $\text{nnz}(\mathbf{A})$  and gap terms using fast subspace embeddings [AC09, CW13] and stochastic optimization [Sha15c]. We obtain significantly improved dependencies on  $\text{sr}(\mathbf{A})$  and  $\epsilon$  and our second running time improves this further when  $\text{nnz}(\mathbf{A}) \leq \frac{d \text{sr}(\mathbf{A})}{\text{gap}^2}$ .
- **Online Eigenvector Estimation:** Given a distribution  $\mathcal{D}$  over vectors  $a \in \mathbb{R}^d$  with covariance matrix  $\Sigma$  and a vector  $x_0$  which is an  $O(\text{gap})$  approximate top eigenvector for  $\Sigma$ , we show how to compute an  $\epsilon$  approximate eigenvector using  $\tilde{O}\left(\frac{v(\mathcal{D})}{\text{gap}^2} + \frac{v(\mathcal{D})}{\text{gap} \cdot \epsilon}\right)$  samples from  $\mathcal{D}$ . Here  $v(\mathcal{D})$  is a natural notion of the variance of  $\mathcal{D}$ . Combining our algorithm with a number of existing algorithms to initialize  $x_0$  we obtain improved sample complexity and runtime results under a variety of assumptions on  $\mathcal{D}$ . Notably, we show that, for general distributions, our sample complexity result is asymptotically optimal - we achieve optimal accuracy as a function of sample size as the number of samples grows large.

We achieve our results using a general framework that we believe is of independent interest. We provide a robust analysis of the classic method of *shift-and-invert* preconditioning to reduce eigenvector computation to *approximately* solving a sequence of linear systems. We then apply variants of stochastic variance reduced gradient descent (SVRG) and additional recent advances in solving linear systems to achieve our claims. We believe our results suggest the generality and effectiveness of shift-and-invert based approaches and imply that further computational improvements may be reaped in practice.

---

\*The manuscript is out of date. An updated version of this work is available at [GHJ<sup>+</sup>16].

# 1 Introduction

Given a matrix  $\mathbf{A} \in \mathbb{R}^{n \times d}$ , computing the top eigenvector of  $\mathbf{A}^\top \mathbf{A}$  is a fundamental problem in computer science with applications ranging from principal component analysis [Jol02], to spectral clustering and learning of mixture models [NJW02, VW04], to pagerank computation [PBMW99], and a number of other graph related computations [Kor03, Spi07].

In this paper we provide improved algorithms for computing the top eigenvector, both in the *offline* case, where the matrix  $\mathbf{A}$  is given explicitly as well as in the *online* or *statistical* case where we are simply given samples from a distribution  $\mathcal{D}$  over vectors  $a \in \mathbb{R}^d$  and wish to compute the top eigenvector of  $\mathbb{E}_{a \sim \mathcal{D}} [aa^\top]$ , the covariance matrix of the distribution.

Our algorithms are based on the classic idea of *shift-and-invert* preconditioning for eigenvalue computation [Saa92]. We give a new robust analysis of the shifted-and-inverted power method, which allows us to efficiently reduce maximum eigenvector computation to *approximately* solving a sequence of linear systems in the matrix  $\lambda \mathbf{I} - \mathbf{A}^\top \mathbf{A}$  for some shift parameter  $\lambda \approx \lambda_1(\mathbf{A})$ . We then show how to solve these systems efficiently using variants of Stochastic Variance Reduced Gradient (SVRG) [JZ13] that optimize a convex function that is given as a sum of non-convex components.

## 1.1 Our Approach

The well known power method for computing the top eigenvector of a  $\mathbf{A}^\top \mathbf{A}$  starts with a initial vector  $x$  (often random) and repeatedly multiplies by  $\mathbf{A}^\top \mathbf{A}$ , eventually causing  $x$  to converge to the top eigenvector. Assuming a random start vector, convergence requires  $O\left(\frac{\log(d/\epsilon)}{\text{gap}}\right)$  iterations, where  $\text{gap} = \frac{\lambda_1 - \lambda_2}{\lambda_1}$ . The dependence on this gap is inherent to the power method and ensures that the largest eigenvalue is significantly amplified in comparison to the remaining values.

If the gap is small, one way to attempt to deal with this dependence is to replace  $\mathbf{A}$  with a preconditioned matrix – i.e. a matrix with the same top eigenvector but a much larger eigenvalue gap. Specifically, let  $\mathbf{B} = \lambda \mathbf{I} - \mathbf{A}$  for some shift parameter  $\lambda$ . We can see that the smallest eigenvector of  $\mathbf{B}$  (the largest eigenvector of  $\mathbf{B}^{-1}$ ) is equal to the largest eigenvector of  $\mathbf{A}$ . Additionally, if  $\lambda$  is near the largest eigenvalue of  $\mathbf{A}$ ,  $\lambda_1$ , there will be a constant gap between the largest and second largest values of  $\mathbf{B}^{-1}$ . For example, if  $\lambda = (1 + \text{gap})\lambda_1$ , then we will have  $\lambda_1(\mathbf{B}^{-1}) = \frac{1}{\lambda - \lambda_1} = \frac{1}{\text{gap} \cdot \lambda_1}$  and  $\lambda_2(\mathbf{B}^{-1}) = \frac{1}{\lambda - \lambda_2} = \frac{1}{2\text{gap} \cdot \lambda_1}$ .

This constant factor gap means that running the power method on  $\mathbf{B}^{-1}$  converges quickly to the top eigenvector of  $\mathbf{A}$ , specifically in  $O(\log(d/\epsilon))$  iterations. Of course, there is a catch – each iteration of the shifted-and-inverted power method requires solving a linear system in  $\mathbf{B}$ . Furthermore, the condition number of  $\mathbf{B}$  is proportional  $\frac{1}{\text{gap}}$ , so as gap gets smaller, solving this linear system becomes more difficult for standard iterative methods.

Fortunately, the problem of solving linear systems is incredibly well studied and there are many efficient iterative algorithms we can adapt to apply  $\mathbf{B}^{-1}$  approximately. In particular, we show how to accelerate the iterations of the shifted-and-inverted power method using variants of Stochastic Variance Reduced Gradient (SVRG) [JZ13]. Due to the condition number of  $\mathbf{B}$ , we will not entirely avoid a  $\frac{1}{\text{gap}}$  dependence, however, we can separate this dependence from the input size  $\text{nnz}(\mathbf{A})$ .

Typically, stochastic gradient methods are used to optimize convex functions that are given as the sum of many convex components. To solve a linear system  $(\mathbf{M}^\top \mathbf{M})x = b$  we minimize the convex function  $f(x) = \frac{1}{2}x^\top (\mathbf{M}^\top \mathbf{M})x + b^\top x$  with components  $\psi_i(x) = \frac{1}{2}x^\top (m_i m_i^\top) x + \frac{1}{n}b^\top x$  where  $m_i$  is the  $i^{\text{th}}$  row of  $\mathbf{M}$ . Such an approach can be used to solve systems in  $\mathbf{A}^\top \mathbf{A}$ , however

solving systems in  $\mathbf{B} = \lambda \mathbf{I} - \mathbf{A}^\top \mathbf{A}$  requires more care. The components of our function are not so simple, and we require an analysis of SVRG that guarantees convergence even when some of these components are *non-convex*. We give a simple analysis for this setting, generalizing recent work appearing in the literature [SS15, CR15].

Given fast approximate solvers for  $\mathbf{B}$ , the second main component required by our algorithmic framework is a new error bound for the shifted-and-inverted power method, showing that it is robust to the approximate linear system solvers, such as SVRG. We give a general robustness analysis, showing exactly what accuracy each system must be solved to, allowing for faster implementations using linear solvers with weaker guarantees. Our proofs center around the potential function

$$G(x) \stackrel{\text{def}}{=} \frac{\|\mathbf{P}_{v_1^\perp} x\|_{\mathbf{B}}}{\|\mathbf{P}_{v_1} x\|_{\mathbf{B}}}$$

where  $\mathbf{P}_{v_1}$  and  $\mathbf{P}_{v_1^\perp}$  are the projections onto the top eigenvector and its complement respectively. This function resembles tangent based potential functions used in previous work [HP14] except that the norms of the projections are measured over  $\mathbf{B}$ . For the exact power method, this is irrelevant – it is not hard to see that progress is identical in both the  $\ell_2$  and  $\mathbf{B}$  norms (see Lemma 33 of the Appendix). However, since  $\|\cdot\|_{\mathbf{B}}$  is a natural norm for measuring the progress of linear system solvers for  $\mathbf{B}$ , our potential function makes it possible to show that progress is also made when we compute  $\mathbf{B}^{-1}x$  approximately up to some error  $\xi$  with bounded  $\|\xi\|_{\mathbf{B}}$ .

## 1.2 Our Results

Our algorithmic framework described above offers several advantages. Theoretically, we obtain improved running times for computing the top eigenvector. In the offline case, in Theorem 16 we give an algorithm running in time  $O\left(\left[\text{nnz}(\mathbf{A}) + \frac{d \text{sr}(\mathbf{A})}{\text{gap}^2}\right] \cdot \left[\log \frac{1}{\epsilon} + \log^2 \frac{d}{\text{gap}}\right]\right)$ , where  $\text{gap}$  is the multiplicative gap between the largest and second largest eigenvalues,  $\text{sr}(\mathbf{A}) = \|\mathbf{A}\|_F^2 / \|\mathbf{A}\|_2^2 \leq \text{rank}(\mathbf{A})$  is the stable rank of  $\mathbf{A}$ , and  $\text{nnz}(\mathbf{A})$  is the number of non-zero entries in the matrix. Up to log factors, our runtime is in many settings proportional to the input size  $\text{nnz}(\mathbf{A})$ , and so is very efficient for large data matrices. In the case when  $\text{nnz}(\mathbf{A}) \leq \frac{d \text{sr}(\mathbf{A})}{\text{gap}^2}$  we also use the results of [FGKS15, LMH15] to provide an accelerated runtime of  $O\left(\left[\frac{\text{nnz}(\mathbf{A})^{3/4} (d \text{sr}(\mathbf{A}))^{1/4}}{\sqrt{\text{gap}}}\right] \cdot \left[\log \frac{d}{\text{gap}} \log \frac{1}{\epsilon} + \log^3 \frac{d}{\text{gap}}\right]\right)$ , shown in Theorem 17.

Our algorithms return an approximate top eigenvector  $x$  with  $x^\top \mathbf{A}^\top \mathbf{A} x \geq (1 - \epsilon) \lambda_1$ . Note that, by choosing error  $\epsilon \cdot \text{gap}$ , we can ensure that  $x$  is actually close to  $v_1$  – i.e. that  $|x^\top v_1| \geq 1 - \epsilon$ . Further, we obtain the same asymptotic runtime since  $O\left(\log \frac{1}{\epsilon \text{gap}} + \log^2 \frac{d}{\text{gap}}\right) = O\left(\log \frac{1}{\epsilon} + \log^2 \frac{d}{\text{gap}}\right)$ . We compare our runtimes with previous work in Table 1.

In the online case, in Theorem 25, we show how to improve an  $O(\text{gap})$  factor approximation to the top eigenvector to an  $\epsilon$  approximation using  $O\left(\frac{v(\mathcal{D}) \log \log(1/\epsilon)}{\text{gap}^2} + \frac{v(\mathcal{D})}{\text{gap} \cdot \epsilon}\right)$  samples where  $v(\mathcal{D})$  is a natural upper bound on the variance of our distribution. Our algorithm is based off the streaming SVRG algorithm of [FGKS14, LMH15]. It requires just  $O(d)$  amortized time to process each sample, uses just  $O(d)$  space, and is easy to parallelize. We can apply our result in a variety of regimes, using algorithms from the literature to obtain the initial  $O(\text{gap})$  factor approximation and our algorithm to refine this solution. As shown in Table 2, this gives improved runtime and sample complexity over existing work. Notably, we give improved asymptotic sample complexity

over known matrix concentration results for general distributions, and give the first streaming algorithm that is asymptotically optimal in the popular Gaussian spike model.

Our bounds hold for *any*  $\mathbf{A}$  or distribution  $\mathcal{D}$ . In the offline case we require no initial knowledge of  $\lambda_1(\mathbf{A})$ , the eigenvalue gap, or the top eigenvector. We are hopeful that our online algorithm can also be made to work without such estimates.

Outside of our runtime results, our robust analysis of the shifted-and-inverted power method provides new understanding of this well studied and widely implemented technique. It gives a means of obtaining provably accurate results when each iteration is implemented using fast approximate linear system solvers with rather weak accuracy guarantees.

In practice, this reduction between approximate linear system solving and eigenvector computation shows that regression libraries can be directly utilized to obtain faster running times for eigenvector computation in many cases. Furthermore, in theory we believe that our reduction suggests computational limits inherent in eigenvector computation as seen by the often easier-to-analyze problem of linear system solving. Indeed, in Section 7, we provide evidence that in certain regimes our statistical results are optimal.

We remark that during the preparation of our manuscript we found that previously and independently Dan Garber and Elad Hazan had discovered a similar technique using shift-and-invert preconditioning and SVRG for sums of non-convex functions to improve the running time for offline eigenvector computation [GH15].

### 1.3 Previous Work

#### Offline Eigenvector Computation

Due to its universal applicability, eigenvector computation in the offline case is extremely well studied. Classical methods, such as the QR algorithm, take roughly  $O(nd^2)$  time to compute a full eigendecomposition. This can be accelerated to  $O(nd^{\omega-1})$ , where  $\omega < 2.373$  is the matrix multiplication constant [Wil12, LG14], however this is still prohibitively expensive for large matrices. Hence, faster iterative methods are often employed, especially when only the top eigenvector (or a few of the top eigenvectors) is desired.

As discussed, the popular power method requires  $O\left(\frac{\log(d/\epsilon)}{\text{gap}}\right)$  iterations to converge to an  $\epsilon$  approximate top eigenvector. Using Chebyshev iteration, or more commonly, the Lanczos method, this bound can be improved to  $O\left(\frac{\log(d/\epsilon)}{\sqrt{\text{gap}}}\right)$  [Saa92], giving total runtime of  $O\left(\text{nnz}(\mathbf{A}) \cdot \frac{\log(d/\epsilon)}{\sqrt{\text{gap}}}\right)$ .

Unfortunately, if  $\mathbf{A}$  is very large and gap is small, this can still be quite expensive, and there is a natural desire to separate the  $\frac{1}{\sqrt{\text{gap}}}$  dependence from the  $\text{nnz}(\mathbf{A})$  term. One approach is to use random subspace embedding matrices [AC09, CW13] or fast row sampling algorithms [CLM<sup>+</sup>15], which can be applied in  $O(\text{nnz}(\mathbf{A}))$  time and yield a matrix  $\tilde{\mathbf{A}}$  which is a good spectral approximation to the original. The number of rows in  $\tilde{\mathbf{A}}$  depends only on the stable rank of  $\mathbf{A}$  and the error of the embedding – hence it can be significantly smaller than  $n$ . Applying such a subspace embedding and then computing the top eigenvector of  $\tilde{\mathbf{A}}$  will require runtime  $O(\text{nnz}(\mathbf{A}) + \text{poly}(\text{sr}(\mathbf{A}), \epsilon, \text{gap}))$ , achieving the goal of reducing runtime dependence on the input size  $\text{nnz}(\mathbf{A})$ . Unfortunately, the dependence on  $\epsilon$  will be significantly suboptimal – such an approach cannot be used to obtain a linearly convergent algorithm. Further, the technique does not extend to online setting, unless we are willing to store a full subspace embedding of our sampled rows.

Another approach, which we follow more closely, is to apply stochastic optimization techniques,

which iteratively update an estimate to the top eigenvector, considering a random row of  $\mathbf{A}$  with each update step. Such algorithms naturally extend to the online setting and have led to improved dependence on the input size for a variety of problems [Bot10]. Using variance-reduced stochastic gradient techniques, [Sha15c] achieves runtime  $O\left(\left(\text{nnz}(\mathbf{A}) + \frac{dr^2n^2}{\text{gap}^2\lambda_1^2}\right) \cdot \log(1/\epsilon) \log \log(1/\epsilon)\right)$  for approximately computing the top eigenvector of a matrix with constant probability. Here  $r$  is an upper bound on the squared row norms of  $\mathbf{A}$ . In the *best case*, when row norms are uniform, this runtime can be simplified to  $O\left(\left(\text{nnz}(\mathbf{A}) + \frac{d\text{sr}(\mathbf{A})^2}{\text{gap}^2}\right) \cdot \log(1/\epsilon) \log \log(1/\epsilon)\right)$ .

The result in [Sha15c] makes an important contribution in separating input size and gap dependencies using stochastic optimization techniques. Unfortunately, the algorithm requires an approximation to the eigenvalue gap and a starting vector that has a constant dot product with the top eigenvector. In [Sha15b] the analysis is extended to a random initialization, however loses polynomial factors in  $d$ . Furthermore, the dependences on the stable rank and  $\epsilon$  are suboptimal – we improve them to  $\text{sr}(\mathbf{A})$  and  $\log(1/\epsilon)$  respectively, obtaining true linear convergence.

Algorithm	Runtime
Power Method	$O\left(\text{nnz}(\mathbf{A}) \frac{\log(d/\epsilon)}{\text{gap}}\right)$
Lanczos Method	$O\left(\text{nnz}(\mathbf{A}) \frac{\log(d/\epsilon)}{\sqrt{\text{gap}}}\right)$
Fast Subspace Embeddings [CW13] Plus Lanczos	$O\left(\text{nnz}(\mathbf{A}) + \frac{d\text{sr}(\mathbf{A})}{\max\{\text{gap}^{2.5}\epsilon, \epsilon^{2.5}\}}\right)$
SVRG [Sha15c] (assuming uniform row norms and warm-start)	$O\left(\left(\text{nnz}(\mathbf{A}) + \frac{d\text{sr}(\mathbf{A})^2}{\text{gap}^2}\right) \cdot \log(1/\epsilon) \log \log(1/\epsilon)\right)$
<b>Theorem 16</b>	$O\left(\left[\text{nnz}(\mathbf{A}) + \frac{d\text{sr}(\mathbf{A})}{\text{gap}^2}\right] \cdot \left[\log \frac{1}{\epsilon} + \log^2 \frac{d}{\text{gap}}\right]\right)$
<b>Theorem 17</b>	$O\left(\left[\frac{\text{nnz}(\mathbf{A})^{3/4}(d\text{sr}(\mathbf{A}))^{1/4}}{\sqrt{\text{gap}}}\right] \cdot \left[\log \frac{d}{\text{gap}} \log \frac{1}{\epsilon} + \log^3 \frac{d}{\text{gap}}\right]\right)$

Table 1: Comparison to Previous work on Offline Eigenvector Estimation. We give runtimes for computing  $x$  such that  $x^\top \mathbf{A}^\top \mathbf{A} x \geq (1 - \epsilon)\lambda_1$ .

## Online Eigenvector Computation

While in the offline case the primary concern is computation time, in the online, or statistical setting, research also focuses on minimizing the number of samples that we must draw from  $\mathcal{D}$  in order to achieve a given accuracy on our eigenvector estimate. Especially sought after are results that achieve asymptotically optimal accuracy as the sample size grows large.

While the result we give in Theorem 25 will work for any distribution parameterized by a variance bound, in this section, in order to more easily compare to previous work, we normalize  $\lambda_1 = 1$  and assume we have the variance bound  $\mathbb{E}_{a \sim \mathcal{D}} \|(aa^\top)^2\|_2 = O(d)$  along with the row norm bound  $\|a\|_2^2 \leq O(d)$ . Additionally, we compare runtimes for computing some  $x$  such that  $|x^\top v_1| \geq 1 - \epsilon$ , as this is the most popular guarantee studied in the literature. Theorem 25 is easily extended to this setting as obtaining  $x$  with  $x^\top \mathbf{A} \mathbf{A}^\top x \geq (1 - \epsilon \cdot \text{gap})\lambda_1$  ensures  $|x^\top v_1| \geq 1 - \epsilon$ . Our algorithm requires  $O\left(\frac{d}{\text{gap}^2\epsilon}\right)$  samples to find such a vector under the assumptions given above.

The simplest algorithm in this setting is to take  $n$  samples from  $\mathcal{D}$  and compute the leading eigenvector of the empirical estimate  $\hat{\mathbb{E}}[aa^\top] = \frac{1}{n} \sum_{i=1}^n a_i a_i^\top$ . By a matrix Bernstein bound, such as

inequality Theorem 6.6.1 of [Tro15],  $O\left(\frac{d \log d}{\text{gap}^2 \epsilon}\right)$  samples is enough to insure  $\left\|\widehat{\mathbb{E}}[aa^\top] - \mathbb{E}[aa^\top]\right\|_2 \leq \sqrt{\epsilon} \text{gap}$ . By Lemma 32 in the Appendix, this gives that, if  $x$  is set to the top eigenvector of  $\widehat{\mathbb{E}}[aa^\top]$  it will satisfy  $|x^\top v_1| \geq 1 - \epsilon$ . Such an  $x$  can then be approximated by applying any offline eigenvector algorithm to the empirical estimate.

A large body of work focuses on improving the computational and sample cost of this simple algorithm, under a variety of assumptions on  $\mathcal{D}$ . The most common focus is on obtaining *streaming algorithms*, in which the storage space is just  $O(d)$  - proportional to the size of a single sample.

In Table 2 we give a sampling of results in this area. All of these results rely on distributional assumptions at least as strong as those given above. In each setting, we can use the cited algorithm to first compute an  $O(\text{gap})$  approximate eigenvector, and then refine this approximation to an  $\epsilon$  approximation using  $O\left(\frac{d}{\text{gap}^2 \epsilon}\right)$  samples by applying our streaming SVRG based algorithm. This allows us to obtain improved runtimes and sample complexities. Notably, by the lower bound shown in Section 7, in all settings considered in Table 2, we achieve optimal asymptotic sample complexity - as our sample size grows large, our  $\epsilon$  decreases at an optimal rate. To save space, we do not include our improved runtime bounds in Table 2, however they are easy to derive by adding the runtime required by the given algorithm to achieve  $O(\text{gap})$  accuracy, to  $O\left(\frac{d^2}{\text{gap}^2 \epsilon}\right)$  - the runtime required by our streaming algorithm.

The bounds given for the simple matrix Bernstein based algorithm described above, Krasulina/Oja's Algorithm [BDF13], and SGD [Sha15a] require no additional assumptions, aside from those given at the beginning of this section. The streaming results cited for [MCJ13] and [HP14] assume  $a$  is generated from a Gaussian spike model, where  $a_i = \sqrt{\lambda_1} \gamma_i v_1 + Z_i$  and  $\gamma_i \sim \mathcal{N}(0, 1)$ ,  $Z_i \sim \mathcal{N}(0, I_d)$ . We note that under this model, the matrix Bernstein results improve by a  $\log d$  factor and so match our results in achieving asymptotically optimal convergence rate. The results of [MCJ13] and [HP14] sacrifice this optimality in order to operate under the streaming model. Our work gives the best of both works - a streaming algorithm giving asymptotically optimal results.

The streaming Aleton algorithm [SRO15] assumes  $\mathbb{E} \|aa^\top \mathbf{W} aa^\top\| \leq O(1) \text{tr}(\mathbf{W})$  for any symmetric  $\mathbf{W}$  that commutes with  $\mathbb{E} aa^\top$ . This is a strictly stronger assumption than the assumption above that  $\mathbb{E}_{a \sim \mathcal{D}} \|(aa^\top)^2\|_2 = O(d)$ .

## 1.4 Paper Organization

**Section 2** Review problem definitions and parameters for our runtime and sample bounds.

**Section 3** Describe the shifted-and-inverted power method and show how it can be implemented using approximate system solvers.

**Section 4** Show how to apply SVRG to solve systems in our shifted matrix, giving our main runtime results for eigenvector computation in the offline setting.

**Section 5** Show how to use an online variant of SVRG to run the shifted-and-inverted power method, giving our main sampling complexity and runtime results in the statistical setting.

**Section 6** Show how to efficiently estimate the shift parameters required by our algorithms, completing their analysis.

**Section 7** Give a lower bound in the statistical setting, showing that our results are asymptotically optimal.

Algorithm	Sample Size	Runtime	Streaming?	Our Sample Complexity
Matrix Bernstein plus Lanczos (explicitly forming sampled matrix)	$O\left(\frac{d \log d}{\text{gap}^2 \epsilon}\right)$	$O\left(\frac{d^3 \log d}{\text{gap}^2 \epsilon}\right)$	$\times$	$O\left(\frac{d \log d}{\text{gap}^3} + \frac{d}{\text{gap}^2 \epsilon}\right)$
Matrix Bernstein plus Lanczos (iteratively applying sampled matrix)	$O\left(\frac{d \log d}{\text{gap}^2 \epsilon}\right)$	$O\left(\frac{d^2 \log d \cdot \log(d/\epsilon)}{\text{gap}^{2.5} \epsilon}\right)$	$\times$	$O\left(\frac{d \log d}{\text{gap}^3} + \frac{d}{\text{gap}^2 \epsilon}\right)$
Memory-efficient PCA [MCJ13, HP14]	$O\left(\frac{d \log(d/\epsilon)}{\text{gap}^3 \epsilon}\right)$	$O\left(\frac{d^2 \log(d/\epsilon)}{\text{gap}^3 \epsilon}\right)$	$\checkmark$	$O\left(\frac{d \log(d/\text{gap})}{\text{gap}^4} + \frac{d}{\text{gap}^2 \epsilon}\right)$
Alecton [SRO15]	$O\left(\frac{d \log(d/\epsilon)}{\text{gap}^2 \epsilon}\right)$	$O\left(\frac{d^2 \log(d/\epsilon)}{\text{gap}^2 \epsilon}\right)$	$\checkmark$	$O\left(\frac{d \log(d/\text{gap})}{\text{gap}^3} + \frac{d}{\text{gap}^2 \epsilon}\right)$
Krasulina / Oja's Algorithm [BDF13]	$O\left(\frac{d^{c_1}}{\text{gap}^2 \epsilon^{c_2}}\right)$	$O\left(\frac{d^{c_1+1}}{\text{gap}^2 \epsilon^{c_2}}\right)$	$\checkmark$	$O\left(\frac{d^{c_1}}{\text{gap}^{2+c_2}} + \frac{d}{\text{gap}^2 \epsilon}\right)$
SGD [Sha15a]	$O\left(\frac{d^3 \log(d/\epsilon)}{\epsilon^2}\right)$	$O\left(\frac{d^4 \log(d/\epsilon)}{\epsilon^2}\right)$	$\checkmark$	$O\left(\frac{d^3 \log(d/\text{gap})}{\text{gap}^2} + \frac{d}{\text{gap}^2 \epsilon}\right)$

Table 2: Summary of existing work on Online Eigenvector Estimation and improvements given by our results. Runtimes are for computing  $x$  such that  $|x^\top v_1| \geq 1 - \epsilon$ . For each of these results we can obtain improved running times and sample complexities by running the algorithm to first compute an  $O(\text{gap})$  approximate eigenvector, and then running our algorithm to obtain an  $\epsilon$  approximation using an additional  $O\left(\frac{d}{\text{gap}^2 \epsilon}\right)$  samples,  $O(d)$  space, and  $O(d)$  work per sample.

## 2 Preliminaries

We bold all matrix variables. We use  $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$ . For a symmetric positive semidefinite (PSD) matrix  $\mathbf{M}$  we let  $\|x\|_{\mathbf{M}} \stackrel{\text{def}}{=} \sqrt{x^\top \mathbf{M} x}$  and we let  $\lambda_1(M), \dots, \lambda_n(M)$  denote its eigenvalues in decreasing order. We use  $\mathbf{M} \preceq \mathbf{N}$  to denote the condition that  $x^\top \mathbf{M} x \leq x^\top \mathbf{N} x$  for all  $x$ .

### 2.1 The Offline Problem

We are given a matrix  $\mathbf{A} \in \mathbb{R}^{n \times d}$  with rows  $a^{(1)}, \dots, a^{(n)} \in \mathbb{R}^d$  and wish to compute an approximation the top eigenvector of  $\Sigma \stackrel{\text{def}}{=} \mathbf{A}^\top \mathbf{A}$ . Specifically for some error parameter  $\epsilon$  we want a unit vector  $x$  such that  $x^\top \Sigma x \geq (1 - \epsilon) \lambda_1(\Sigma)$ .

### 2.2 The Statistical Problem

We are given  $n$  independent samples from a distribution  $\mathcal{D}$  on  $\mathbb{R}^d$  and wish to compute the top eigenvector of  $\Sigma \stackrel{\text{def}}{=} \mathbb{E}_{a \sim \mathcal{D}} [aa^\top]$ . Again, for some error parameter  $\epsilon$  we want to return a unit vector  $x$  such that  $x^\top \Sigma x \geq (1 - \epsilon) \lambda_1(\Sigma)$ .

### 2.3 Problem Parameters

We parameterize the running time of our algorithm in terms of several natural properties of  $\mathbf{A}$ ,  $\mathcal{D}$ , and  $\Sigma$ . We let  $\lambda_1, \dots, \lambda_d$  denote the eigenvalues of  $\Sigma$  in decreasing order and we let  $v_1, \dots, v_d$  denote their corresponding eigenvectors. We define the *eigenvalue gap* by  $\text{gap} \stackrel{\text{def}}{=} \frac{\lambda_1 - \lambda_2}{\lambda_1}$ .

We use the following additional parameters to provide running times for the offline and statistical problems respectively:

- **Offline Problem:** We let  $\text{sr}(\mathbf{A}) \stackrel{\text{def}}{=} \sum_i \frac{\lambda_i}{\lambda_1} = \frac{\|\mathbf{A}\|_F^2}{\|\mathbf{A}\|_2^2}$  denote the stable rank of  $\mathbf{A}$ . Note that we always have  $\text{sr}(\mathbf{A}) \leq \text{rank}(\mathbf{A})$ . We let  $\text{nnz}(\mathbf{A})$  denote the number of non-zero entries in  $\mathbf{A}$ .
- **Online Problem:** We let  $v(\mathcal{D}) \stackrel{\text{def}}{=} \frac{\|\mathbb{E}_{a \sim \mathcal{D}}[(aa^\top)^2]\|_2}{\|\mathbb{E}_{a \sim \mathcal{D}}(aa^\top)\|_2^2} = \frac{\|\mathbb{E}_{a \sim \mathcal{D}}[(aa^\top)^2]\|_2}{\lambda_1^2}$  denote a natural upper bound on the variance of  $\mathcal{D}$  in various settings. Note that  $v(\mathcal{D}) \geq 1$ .

## 3 Framework

Here we develop our robust shift-and-invert framework. In Section 3.1 we provide a basic overview of the framework and in Section 3.2 we introduce the potential function we use to measure progress of our algorithms. In Section 3.3 we show how to analyze the framework given access to an exact linear system solver and in Section 3.4 we strengthen this analysis to work with an inexact linear system solver. Finally, in Section 3.5 we discuss initializing the framework.

### 3.1 Shifted-and-Inverted Power Method Basics

We let  $\mathbf{B}_\lambda \stackrel{\text{def}}{=} \lambda \mathbf{I} - \Sigma$  denote the shifted matrix that we will use in our implementation of the shifted-and-inverted power method. As discussed, in order for  $\mathbf{B}_\lambda^{-1}$  to have a large eigenvalue gap,  $\lambda$  should be set to  $(1 + c \cdot \text{gap}) \lambda_1$  for some constant  $c \geq 0$ . Throughout this section we assume that we have



a crude estimate of  $\lambda_1$  and  $\text{gap}$  and fix  $\lambda$  to be a value satisfying  $(1 + \frac{\text{gap}}{150}) \lambda_1 \leq \lambda \leq (1 + \frac{\text{gap}}{100}) \lambda_1$ . (See Section 6 for how we can compute such a  $\lambda$ ). For the remainder of this section we work with such a fixed value of  $\lambda$  and therefore for convenience denote  $\mathbf{B}_\lambda$  as  $\mathbf{B}$ .

Note that  $\lambda_i(\mathbf{B}^{-1}) = \frac{1}{\lambda_i(\mathbf{B})} = \frac{1}{\lambda - \lambda_i}$  and so  $\frac{\lambda_1(\mathbf{B}^{-1})}{\lambda_2(\mathbf{B}^{-1})} = \frac{\lambda - \lambda_2}{\lambda - \lambda_1} \geq \frac{\text{gap}}{\text{gap}/100} = 100$ . This large gap will ensure that, assuming the ability to apply  $\mathbf{B}^{-1}$ , the power method will converge very quickly. In the remainder of this section we develop our error analysis for the shifted-and-inverted power method which demonstrates that approximate application of  $\mathbf{B}^{-1}$  in each iteration in fact suffices.

### 3.2 Potential Function

Our analysis of the power method focuses on the objective of maximizing the Rayleigh quotient,  $x^\top \Sigma x$  for unit vector  $x$ . Note that as the following lemma shows, this has a directly correspondence to the error in maximizing  $|v_1^\top x|$ :

**Lemma 1** (Bounding Eigenvector Error by Rayleigh Quotient). *For a unit vector  $x$  let  $\epsilon = \lambda_1 - x^\top \Sigma x$ . If  $\epsilon \leq \lambda_1 \cdot \text{gap}$  then*

$$|v_1^\top x| \geq \sqrt{1 - \frac{\epsilon}{\lambda_1 \cdot \text{gap}}}.$$

*Proof.* Among all unit vectors  $x$  such that  $\epsilon = \lambda_1 - x^\top \Sigma x$ , a minimizer of  $|v_1^\top x|$  has the form  $x = (\sqrt{1 - \delta^2})v_1 + \delta v_2$  for some  $\delta$ . We have

$$\epsilon = \lambda_1 - x^\top \Sigma x = \lambda_1 - \lambda_1(1 - \delta^2) - \lambda_2\delta^2 = (\lambda_1 - \lambda_2)\delta^2.$$

Therefore by direct computation,

$$|v_1^\top x| = \sqrt{1 - \delta^2} = \sqrt{1 - \frac{\epsilon}{\lambda_1 - \lambda_2}} = \sqrt{1 - \frac{\epsilon}{\lambda_1 \cdot \text{gap}}}.$$

□

In order to track the progress of our algorithm we use a more complex potential function than just the Rayleigh quotient error,  $\lambda_1 - x^\top \Sigma x$ . Our potential function  $G$  is defined for  $x \neq 0$  by

$$G(x) \stackrel{\text{def}}{=} \frac{\|\mathbf{P}_{v_1^\perp}(x)\|_{\mathbf{B}}}{\|\mathbf{P}_{v_1}(x)\|_{\mathbf{B}}}$$

where  $\mathbf{P}_{v_1}(x)$  and  $\mathbf{P}_{v_1^\perp}(x)$  denote the projections of  $x$  in the direction of  $v_1$  and on the subspace orthogonal to  $v_1$  respectively. Equivalently, we have that:

$$G(x) = \frac{\sqrt{\|x\|_{\mathbf{B}}^2 - (v_1^\top \mathbf{B}^{1/2} x)^2}}{|v_1^\top \mathbf{B}^{1/2} x|} = \frac{\sqrt{\sum_{i \geq 2} \frac{\alpha_i^2}{\lambda_i(\mathbf{B}^{-1})}}}{\sqrt{\frac{\alpha_1^2}{\lambda_1(\mathbf{B}^{-1})}}}. \quad (1)$$

where  $\alpha_i = v_i^\top x$ .

When the Rayleigh quotient error  $\epsilon = \lambda_1 - x^\top \Sigma x$  of  $x$  is small, we can show a strong relation between  $\epsilon$  and  $G(x)$ . We prove this in two parts. First we prove a technical lemma, Lemma 2, that we will use several times for bounding the numerator of  $G$  and then we prove the connection in Lemma 3.

**Lemma 2.** For a unit vector  $x$  and  $\epsilon = \lambda_1 - x^\top \Sigma x$  if  $\epsilon \leq \lambda_1 \cdot \text{gap}$  then

$$\epsilon \leq x^\top \mathbf{B}x - (v_1^\top \mathbf{B}x)(v_1^\top x) \leq \epsilon \left(1 + \frac{\lambda - \lambda_1}{\lambda_1 \cdot \text{gap}}\right).$$

*Proof.* Since  $\mathbf{B} = \lambda \mathbf{I} - \Sigma$  and since  $v_1$  is an eigenvector of  $\Sigma$  with eigenvalue  $\lambda_1$  we have

$$\begin{aligned} x^\top \mathbf{B}x - (v_1^\top \mathbf{B}x)(v_1^\top x) &= \lambda \|x\|_2^2 - x^\top \Sigma x - (\lambda v_1^\top x - v_1^\top \Sigma x)(v_1^\top x) \\ &= \lambda - \lambda_1 + \epsilon - (\lambda v_1^\top x - \lambda_1 v_1^\top x)(v_1^\top x) \\ &= (\lambda - \lambda_1) \left(1 - (v_1^\top x)^2\right) + \epsilon. \end{aligned}$$

Now by Lemma 1 we know that  $|v_1^\top x| \geq \sqrt{1 - \frac{\epsilon}{\lambda_1 \cdot \text{gap}}}$ , giving us the upper bound. Furthermore, since trivially  $|v_1^\top x| \leq 1$  and  $\lambda - \lambda_1 > 0$ , we have the lower bound.  $\square$

**Lemma 3** (Potential Function to Rayleigh Quotient Error Conversion). For a unit vector  $x$  and  $\epsilon = \lambda_1 - x^\top \Sigma x$  if  $\epsilon \leq \frac{1}{2} \lambda_1 \cdot \text{gap}$ , we have:

$$\frac{\epsilon}{\lambda - \lambda_1} \leq G(x)^2 \leq \left(1 + \frac{\lambda - \lambda_1}{\lambda_1 \cdot \text{gap}}\right) \left(1 + \frac{2\epsilon}{\lambda_1 \cdot \text{gap}}\right) \frac{\epsilon}{\lambda - \lambda_1}.$$

*Proof.* Since  $v_1$  is an eigenvector of  $\mathbf{B}$ , we can write  $G(x)^2 = \frac{x^\top \mathbf{B}x - (v_1^\top \mathbf{B}x)(v_1^\top x)}{(v_1^\top \mathbf{B}x)(v_1^\top x)}$ . Lemmas 1 and 2 then give us:

$$\frac{\epsilon}{\lambda - \lambda_1} \leq G(x)^2 \leq \left(1 + \frac{\lambda - \lambda_1}{\lambda_1 \cdot \text{gap}}\right) \frac{\epsilon}{(\lambda - \lambda_1) \left(1 - \frac{\epsilon}{\lambda_1 \cdot \text{gap}}\right)}.$$

Since  $\epsilon \leq \frac{1}{2} \lambda_1 \cdot \text{gap}$ , we have  $\frac{1}{1 - \frac{\epsilon}{\lambda_1 \cdot \text{gap}}} \leq 1 + \frac{2\epsilon}{\lambda_1 \cdot \text{gap}}$ . This proves the lemma.  $\square$

### 3.3 Power Iteration

Here we show that the shifted-and-inverted power iteration in fact makes progress with respect to our objective function given an exact linear system solver for  $\mathbf{B}$ . Formally, we show that applying  $\mathbf{B}^{-1}$  to a vector  $x$  decreases the potential function  $G(x)$  geometrically.

**Theorem 4.** Let  $x$  be a unit vector with  $\langle x, v_1 \rangle \neq 0$  and let  $\tilde{x} = \mathbf{B}^{-1}x$ , i.e. the power method update of  $\mathbf{B}^{-1}$  on  $x$ . Then, we have:

$$G(\tilde{x}) \leq \frac{\lambda_2(\mathbf{B}^{-1})}{\lambda_1(\mathbf{B}^{-1})} G(x) \leq \frac{1}{100} G(x).$$

Note that  $\tilde{x}$  may no longer be a unit vector. However,  $G(\tilde{x}, v_1) = G(c\tilde{x}, v_1)$  for any scaling parameter  $c$ , so the theorem also holds for  $\tilde{x}$  scaled to have unit norm.

*Proof.* Writing  $x$  in the eigenbasis, we have  $x = \sum_i \alpha_i v_i$  and  $\tilde{x} = \sum_i \alpha_i \lambda_i (\mathbf{B}^{-1}) v_i$ . Since  $\langle x, v_1 \rangle \neq 0$ ,  $\alpha_1 \neq 0$  and by the equivalent formation of  $G(x)$  given in (1):

$$G(\tilde{x}) = \frac{\sqrt{\sum_{i \geq 2} \alpha_i^2 \lambda_i(\mathbf{B}^{-1})}}{\sqrt{\alpha_1^2 \lambda_1(\mathbf{B}^{-1})}} \leq \frac{\lambda_2(\mathbf{B}^{-1})}{\lambda_1(\mathbf{B}^{-1})} \cdot \frac{\sqrt{\sum_{i \geq 2} \frac{\alpha_i^2}{\lambda_i(\mathbf{B}^{-1})}}}{\sqrt{\frac{\alpha_1^2}{\lambda_1(\mathbf{B}^{-1})}}} = \frac{\lambda_2(\mathbf{B}^{-1})}{\lambda_1(\mathbf{B}^{-1})} \cdot G(x).$$

Recalling that  $\frac{\lambda_1(\mathbf{B}^{-1})}{\lambda_2(\mathbf{B}^{-1})} = \frac{\lambda - \lambda_2}{\lambda - \lambda_1} \geq \frac{\text{gap}}{\text{gap}/100} = 100$  yields the result.  $\square$

The challenge in using the above theorem, and any traditional analysis of the shifted-and-inverted power method, is that we don't actually have access to  $\mathbf{B}^{-1}$ . In the next section we show that the shifted-and-inverted power method is robust – we still make progress on our objective function even if we only approximate  $\mathbf{B}^{-1}x$  using a fast linear system solver.

### 3.4 Approximate Power Iteration

We are now ready to prove our main result on the shifted-and-inverted power method using approximate linear system solves at each iteration. In words, we show that each iteration of the power method makes constant factor expected progress on our potential function assuming we:

1. Start with a sufficiently good  $x$  and an approximation of  $\lambda_1$
2. Apply  $\mathbf{B}^{-1}$  approximately using a system solver such that the function error (or distance to  $\mathbf{B}^{-1}x$  in the  $\mathbf{B}$  norm) is sufficiently small *in expectation*.
3. Estimate Rayleigh quotients over  $\Sigma$  well enough to only accept updates that do not hurt progress on the objective function too much.

This third assumption is necessary since the second assumption is quite weak. An expected progress bound on the linear system solver allows, for example, the solver to occasionally return a solution that is entirely orthogonal to  $v_1$ , causing us to make unbounded backwards progress on our potential function. The third assumption allows us to reject possibly harmful updates and ensure that we still make progress in expectation. In the offline setting, we can access  $\mathbf{A}$  and are able to compute Rayleigh quotients exactly in time  $\text{nnz}(\mathbf{A})$  time. However, we only assume the ability to estimate quotients as in the online setting we only have access to  $\Sigma$  through samples from  $\mathcal{D}$ .

While, our general theorem for the approximate power iteration, Theorem 5, assumes that we can solve linear systems to some absolute accuracy in expectation, this is not the standard assumption for many linear system solvers. Many fast iterative linear system solvers assume an initial approximation to  $\mathbf{B}^{-1}x$  and then show that the quality of this approximation is then improved geometrically in each iteration of the algorithm. In Corollary 6 we show how to find a coarse initial approximation to  $\mathbf{B}^{-1}x$ , in fact just approximating  $\mathbf{B}^{-1}$  with  $\frac{1}{x^\top \mathbf{B}x}x$ . Moreover, using this coarse approximation in Corollary 6 we show that Theorem 5 actually implies that it suffices to just make a fixed relative amount of progress in solving the linear system.

Note that in both claims we measure error of the linear system solver using  $\|\cdot\|_{\mathbf{B}}$ . This is a natural norm in which geometric convergence is shown for many linear system solvers and directly corresponds to the function error of minimizing  $f(w) = \frac{1}{2}w^\top \mathbf{B}w - w^\top x$  to compute  $\mathbf{B}^{-1}x$ .

**Theorem 5** (Approximate Shifted-and-Inverted Power Iteration – Warm Start). *Let  $x = \sum_i \alpha_i v_i$  be a unit vector such that  $G(x) \leq \frac{1}{\sqrt{10}}$ . Suppose we know some shift parameter  $\lambda$  with  $(1 + \frac{\text{gap}}{150}) \lambda_1 < \lambda \leq (1 + \frac{\text{gap}}{100}) \lambda_1$  and an estimate  $\hat{\lambda}_1$  of  $\lambda_1$  such that  $\frac{10}{11}(\lambda - \lambda_1) \leq \lambda - \hat{\lambda}_1 \leq \lambda - \lambda_1$ . Furthermore, suppose we have a subroutine  $\text{solve}(\cdot)$  that on any input  $x$*

$$\mathbb{E} [\|\text{solve}(x) - \mathbf{B}^{-1}x\|_{\mathbf{B}}] \leq \frac{c_1}{1000} \sqrt{\lambda_1(\mathbf{B}^{-1})},$$

for some  $c_1 < 1$ , and a subroutine  $\widehat{\text{quot}}(\cdot)$  that on any input  $x \neq 0$

$$\left| \widehat{\text{quot}}(x) - \text{quot}(x) \right| \leq \frac{1}{30} (\lambda - \lambda_1) \quad \text{for all nonzero } x \in \mathbb{R}^d.$$

where  $\text{quot}(x) \stackrel{\text{def}}{=} \frac{x^\top \Sigma x}{x^\top x}$ .

Then the following update procedure:

$$\text{Set } \hat{x} = \text{solve}(x),$$

$$\text{Set } \tilde{x} = \begin{cases} \hat{x} & \text{if } \left\{ \begin{array}{l} \widehat{\text{quot}}(\hat{x}) \geq \hat{\lambda}_1 - (\lambda - \hat{\lambda}_1)/6 \text{ and} \\ \|\hat{x}\|_2 \geq \frac{2}{3} \frac{1}{\lambda - \hat{\lambda}_1} \end{array} \right. \\ x & \text{otherwise,} \end{cases}$$

satisfies the following:

- $G(\tilde{x}) \leq \frac{1}{\sqrt{10}}$  and
- $\mathbb{E}[G(\tilde{x})] \leq \frac{3}{25} G(x) + \frac{c_1}{500}$ .

That is, not only do we decrease our potential function by a constant factor in expectation, but we are guaranteed that the potential function will never increase beyond  $1/\sqrt{10}$ .

*Proof.* The first claim follows directly from our choice of  $\tilde{x}$  from  $x$  and  $\hat{x}$ . If  $\tilde{x} = x$ , it holds trivially by our assumption that  $G(x) \leq \frac{1}{\sqrt{10}}$ . Otherwise,  $\tilde{x} = \hat{x}$  and we know that

$$\begin{aligned} \lambda_1 - \text{quot}(\hat{x}) &\leq \hat{\lambda}_1 - \text{quot}(\hat{x}) \leq \hat{\lambda}_1 - \widehat{\text{quot}}(\hat{x}) + \left| \widehat{\text{quot}}(\hat{x}) - \text{quot}(\hat{x}) \right| \\ &\leq \frac{\lambda - \hat{\lambda}_1}{6} + \frac{\lambda - \lambda_1}{30} \leq \frac{\lambda - \lambda_1}{5} \leq \frac{\lambda_1 \cdot \text{gap}}{500}. \end{aligned}$$

The claim then follows from Lemma 3 as

$$\begin{aligned} G(\hat{x})^2 &\leq \left( 1 + \frac{\lambda - \lambda_1}{\lambda_1 \cdot \text{gap}} \right) \left( 1 + \frac{2(\lambda_1 - \text{quot}(\hat{x}))}{\lambda_1 \cdot \text{gap}} \right) \frac{\lambda_1 - \text{quot}(\hat{x})}{\lambda - \lambda_1} \\ &\leq \frac{101}{100} \cdot \frac{251}{250} \cdot \frac{\left( \frac{\lambda_1 \cdot \text{gap}}{500} \right)}{\left( \frac{\lambda_1 \cdot \text{gap}}{150} \right)} \leq \frac{1}{\sqrt{10}}. \end{aligned}$$

All that remains is to show the second claim, that  $\mathbb{E}[G(\tilde{x})] \leq \frac{3}{25} G(x) + \frac{4c_1}{1000}$ . Let  $\mathcal{F}$  denote the event that we accept our iteration and set  $x = \hat{x} = \text{solve}(x)$ . That is:

$$\mathcal{F} \stackrel{\text{def}}{=} \left\{ \widehat{\text{quot}}(\hat{x}) \geq \hat{\lambda}_1 - \frac{\lambda - \hat{\lambda}_1}{6} \right\} \cup \left\{ \|\hat{x}\|_2 \geq \frac{2}{3} \frac{1}{\lambda - \hat{\lambda}_1} \right\}.$$

Using our bounds on  $\hat{\lambda}_1$  and  $\widehat{\text{quot}}(\cdot)$ , we know that  $\widehat{\text{quot}}(x) \leq \text{quot}(x) + (\lambda - \lambda_1)/30$  and  $\lambda - \hat{\lambda}_1 \leq \lambda - \lambda_1$ . Therefore, since  $-1/6 - 1/30 \geq -1/2$  we have

$$\mathcal{F} \subseteq \{ \text{quot}(\hat{x}) \geq \lambda_1 - (\lambda - \lambda_1)/2 \} \cup \left\{ \|\hat{x}\|_2 \geq \frac{2}{3} \frac{1}{\lambda - \lambda_1} \right\},$$

We will complete the proof in two steps. First we let  $\xi \stackrel{\text{def}}{=} \hat{x} - \mathbf{B}^{-1}x$  and show that assuming  $\mathcal{F}$  is true then  $G(\hat{x})$  and  $\|\xi\|_{\mathbf{B}}$  are linearly related, i.e. expected error bounds on  $\|\xi\|_{\mathbf{B}}$  correspond to expected error bounds on  $G(\hat{x})$ . Second, we bound the probability that  $\mathcal{F}$  does not occur and bound error incurred in this case. Combining these yields the result.

To show the linear relationship in the case where  $\mathcal{F}$  is true, first note Lemma 1 shows that in this case  $\left|v_1^\top \frac{\hat{x}}{\|\hat{x}\|_2}\right| \geq \sqrt{1 - \frac{\lambda_1 - \text{quot}(\hat{x})}{\lambda_1 \cdot \text{gap}}} \geq \frac{3}{4}$ . Consequently,

$$\|\mathbf{P}_{v_1}(\hat{x})\|_{\mathbf{B}} = \left|v_1^\top \hat{x}\right| \sqrt{\lambda - \lambda_1} = \left|v_1^\top \frac{\hat{x}}{\|\hat{x}\|_2}\right| \cdot \|\hat{x}\| \sqrt{\lambda - \lambda_1} \geq \frac{3}{4} \cdot \frac{2}{3} \frac{1}{\sqrt{\lambda - \lambda_1}} = \frac{\sqrt{\lambda_1(\mathbf{B}^{-1})}}{2}.$$

However,

$$\left\|\mathbf{P}_{v_1^\perp}(\hat{x})\right\|_{\mathbf{B}} \leq \left\|\mathbf{P}_{v_1^\perp}(\mathbf{B}^{-1}x)\right\|_{\mathbf{B}} + \left\|\mathbf{P}_{v_1^\perp}(\xi)\right\|_{\mathbf{B}} \leq \left\|\mathbf{P}_{v_1^\perp}(\mathbf{B}^{-1}x)\right\|_{\mathbf{B}} + \|\xi\|_{\mathbf{B}}$$

and by Theorem 4 and the definition of  $G$  we have

$$\left\|\mathbf{P}_{v_1^\perp}(\mathbf{B}^{-1}x)\right\|_{\mathbf{B}} = \left\|\mathbf{P}_{v_1}(\mathbf{B}^{-1}x)\right\|_{\mathbf{B}} \cdot G(\mathbf{B}^{-1}x) \leq \left(|\langle x, v_1 \rangle| \sqrt{\lambda_1(\mathbf{B}^{-1})}\right) \cdot \frac{G(x)}{100}.$$

Taking expectations, using that  $|\langle x, v_1 \rangle| \leq 1$ , and combining these three inequalities yields

$$\mathbb{E}[G(\hat{x})|\mathcal{F}] = \mathbb{E}\left[\frac{\left\|\mathbf{P}_{v_1^\perp}(\mathbf{B}^{-1}x)\right\|_{\mathbf{B}}}{\left\|\mathbf{P}_{v_1}(\mathbf{B}^{-1}x)\right\|_{\mathbf{B}}}\middle|\mathcal{F}\right] \leq \frac{G(x)}{50} + 2 \frac{\mathbb{E}[\|\xi\|_{\mathbf{B}}|\mathcal{F}]}{\sqrt{\lambda_1(\mathbf{B}^{-1})}} \quad (2)$$

So, conditioning on making an update and changing  $x$  (i.e.  $\mathcal{F}$  occurring), we see that our potential function changes exactly as in the exact case (Theorem 4) with additional additive error due to our inexact linear system solve.

Next we upper bound  $\mathbb{P}[\mathcal{F}]$  and use it to compute  $\mathbb{E}[\|\xi\|_{\mathbf{B}}|\mathcal{F}]$ . We will show that

$$\mathcal{G} \stackrel{\text{def}}{=} \left\{\|\xi\|_{\mathbf{B}} \leq \frac{1}{100} \cdot \sqrt{\lambda_1(\mathbf{B}^{-1})}\right\} \subseteq \mathcal{F}$$

which then implies by Markov inequality that

$$\mathbb{P}[\mathcal{F}] \geq \mathbb{P}\left[\|\xi\|_{\mathbf{B}} \leq \frac{1}{100} \cdot \sqrt{\lambda_1(\mathbf{B}^{-1})}\right] \geq 1 - \frac{\mathbb{E}[\|\xi\|_{\mathbf{B}}]}{\frac{1}{100} \cdot \sqrt{\lambda_1(\mathbf{B}^{-1})}} \geq \frac{9}{10}, \quad (3)$$

where we used the fact that  $\mathbb{E}[\|\xi\|_{\mathbf{B}}] \leq \frac{c_1}{1000} \sqrt{\lambda_1(\mathbf{B}^{-1})}$  for some  $c_1 < 1$ .

Let us now show that  $\mathcal{G} \subseteq \mathcal{F}$ . Suppose  $\mathcal{G}$  is occurs. We can bound  $\|\hat{x}\|_2$  as follows:

$$\begin{aligned} \|\hat{x}\|_2 &\geq \|\mathbf{B}^{-1}x\|_2 - \|\xi\|_2 \geq \|\mathbf{B}^{-1}x\| - \sqrt{\lambda_1(\mathbf{B}^{-1})} \|\xi\|_{\mathbf{B}} \\ &\geq |\alpha_1| \lambda_1(\mathbf{B}^{-1}) - \frac{1}{100} \cdot \lambda_1(\mathbf{B}^{-1}) \\ &= \frac{1}{\lambda - \lambda_1} \left(|\alpha_1| - \frac{1}{100}\right) \geq \frac{3}{4} \frac{1}{\lambda - \lambda_1}, \end{aligned} \quad (4)$$

where we use Lemmas 2 and 3 to conclude that  $|\alpha_1| \geq \sqrt{1 - \frac{1}{10}}$ . We now turn to showing the Rayleigh quotient condition required by  $\mathcal{F}$ . In order to do this, we first bound  $\hat{x}^\top \mathbf{B} \hat{x} - (v_1^\top \mathbf{B} \hat{x}) (v_1^\top \hat{x})$  and then use Lemma 2. We have:

$$\begin{aligned} \sqrt{\hat{x}^\top \mathbf{B} \hat{x} - (v_1^\top \mathbf{B} \hat{x}) (v_1^\top \hat{x})} &= \left\| \mathbf{P}_{v_1^\perp}(\hat{x}) \right\|_{\mathbf{B}} \leq \left\| \mathbf{P}_{v_1^\perp}(\mathbf{B}^{-1}x) \right\|_{\mathbf{B}} + \left\| \mathbf{P}_{v_1^\perp}(\xi) \right\|_{\mathbf{B}} \\ &\leq \sqrt{\sum_{i \geq 2} \alpha_i^2 \lambda_i (\mathbf{B}^{-1})} + \frac{1}{100} \cdot \sqrt{\lambda_1 (\mathbf{B}^{-1})} \\ &\leq \sqrt{\lambda_2 (\mathbf{B}^{-1})} + \frac{1}{100} \cdot \sqrt{\lambda_1 (\mathbf{B}^{-1})} \leq \frac{1}{9} \sqrt{\lambda - \lambda_1}, \end{aligned}$$

where we used the fact that  $\lambda_2 (\mathbf{B}^{-1}) \leq \frac{1}{100} \lambda_1 (\mathbf{B}^{-1})$  since  $\lambda \leq \lambda_1 + \frac{\text{gap}}{100}$  in the last step. Now, using Lemma 2 and the bound on  $\|\hat{x}\|_2$ , we conclude that

$$\begin{aligned} \hat{\lambda}_1 - \widehat{\text{quot}}(\hat{x}) &\leq \lambda_1 - \text{quot}(\hat{x}) + \left| \text{quot}(\hat{x}) - \widehat{\text{quot}}(\hat{x}) \right| + \hat{\lambda}_1 - \lambda_1 \\ &\leq \frac{\hat{x}^\top \mathbf{B} \hat{x} - (v_1^\top \mathbf{B} \hat{x}) (v_1^\top \hat{x})}{\|\hat{x}\|_2^2} + \frac{\lambda - \lambda_1}{30} + \frac{\lambda - \lambda_1}{11} \\ &\leq \frac{1}{81(\lambda - \lambda_1)} \cdot \frac{16}{9} (\lambda - \lambda_1)^2 + \frac{\lambda - \lambda_1}{8} \\ &\leq (\lambda - \lambda_1) / 6 \leq (\lambda - \hat{\lambda}_1) / 4. \end{aligned} \tag{5}$$

Combining (4) and (5) shows that  $\mathcal{G} \subseteq \mathcal{F}$  there by proving (3).

Using this and the fact that  $\|\cdot\|_{\mathbf{B}} \geq 0$  we can upper bound  $\mathbb{E}[\|\xi\|_{\mathbf{B}}|\mathcal{F}]$  as follows:

$$\mathbb{E}[\|\xi\|_{\mathbf{B}}|\mathcal{F}] \leq \frac{1}{\mathbb{P}[\mathcal{F}]} \cdot \mathbb{E}[\|\xi\|_{\mathbf{B}}] \leq \frac{c_1}{900} \cdot \sqrt{\lambda_1 (\mathbf{B}^{-1})}$$

Plugging this into (2), we obtain:

$$\mathbb{E}[G(\hat{x})|\mathcal{F}] \leq \frac{1}{50} G(x) + \frac{2\mathbb{E}[\|\xi\|_{\mathbf{B}}|\mathcal{F}]}{\sqrt{\lambda_1 (\mathbf{B}^{-1})}} \leq \frac{1}{50} \cdot G(x) + \frac{2c_1}{900}.$$

We can now finally bound  $\mathbb{E}[G(\tilde{x})]$  as follows:

$$\begin{aligned} \mathbb{E}[G(\tilde{x})] &= \mathbb{P}[\mathcal{F}] \cdot \mathbb{E}[G(\hat{x})|\mathcal{F}] + (1 - \mathbb{P}[\mathcal{F}]) G(x) \\ &\leq \frac{9}{10} \left( \frac{1}{50} \cdot G(x) + \frac{2c_1}{900} \right) + \frac{1}{10} G(x) = \frac{3}{25} G(x) + \frac{2c_1}{1000}. \end{aligned}$$

This proves the theorem. □

**Corollary 6** (Relative Error Linear System Solvers). *For any unit vector  $x$ , we have:*

$$\left\| \frac{1}{x^\top \mathbf{B} x} x - \mathbf{B}^{-1} x \right\|_{\mathbf{B}} \leq \alpha_1 \sqrt{\lambda_1 (\mathbf{B}^{-1})} \cdot G(x) = \lambda_1 (\mathbf{B}^{-1}) \sqrt{\sum_{i \geq 2} \frac{\alpha_i^2}{\lambda_i (\mathbf{B}^{-1})}}, \tag{6}$$

where  $x = \sum_i \alpha_i v_i$  is the decomposition of  $x$  along  $v_i$ . Therefore, instantiating Theorem 5 with  $c_1 = \alpha_1 G(x)$  gives  $\mathbb{E}[G(\tilde{x})] \leq \frac{4}{25} G(x)$  as long as:

$$\mathbb{E} [\|\text{solve}(x) - \mathbf{B}^{-1}x\|_{\mathbf{B}}] \leq \frac{1}{1000} \left\| \frac{1}{\lambda - x^\top \Sigma x} x - \mathbf{B}^{-1}x \right\|_{\mathbf{B}}.$$

*Proof.* Since  $\mathbf{B}$  is PSD we see that if we let  $f(w) = \frac{1}{2} w^\top \mathbf{B} w - w^\top x$ , then the minimizer is  $\mathbf{B}^{-1}x$ . Furthermore note that  $\frac{1}{x^\top \mathbf{B} x} = \arg \min_{\beta} f(\beta x)$  and therefore

$$\begin{aligned} \left\| \frac{1}{x^\top \mathbf{B} x} x - \mathbf{B}^{-1}x \right\|_{\mathbf{B}}^2 &= x^\top \mathbf{B}^{-1}x - \frac{1}{x^\top \mathbf{B} x} = 2 \left[ f\left(\frac{x}{x^\top \mathbf{B} x}\right) - f(\mathbf{B}^{-1}x) \right] \\ &= 2 \left[ \min_{\beta} f(\beta x) - f(\mathbf{B}^{-1}x) \right] \leq 2 \left[ f(\lambda_1(\mathbf{B}^{-1})x) - f(\mathbf{B}^{-1}x) \right] \\ &= \lambda_1(\mathbf{B}^{-1})^2 x^\top \mathbf{B} x - 2\lambda_1(\mathbf{B}^{-1}) x^\top x + x^\top \mathbf{B}^{-1}x \\ &= \sum_{i=1}^d \left| v_i^\top \mathbf{B}^{\frac{1}{2}} x \right|^2 (\lambda_1(\mathbf{B}^{-1}) - \lambda_i(\mathbf{B}^{-1}))^2 \leq \lambda_1(\mathbf{B}^{-1})^2 \sum_{i \geq 2} \left| v_i^\top \mathbf{B}^{\frac{1}{2}} x \right|^2 \\ &= \lambda_1(\mathbf{B}^{-1})^2 \sum_{i \geq 2} \frac{\alpha_i^2}{\lambda_i(\mathbf{B}^{-1})}, \end{aligned}$$

which proves (6).

Consequently

$$\frac{c_1}{1000} \sqrt{\lambda_1(\mathbf{B}^{-1})} = \frac{1}{1000} \alpha_1 G(x) \sqrt{\lambda_1(\mathbf{B}^{-1})} \geq \frac{1}{1000} \left\| \frac{1}{x^\top \mathbf{B} x} x - \mathbf{B}^{-1}x \right\|_{\mathbf{B}}$$

which with Theorem 5 then completes the proof.  $\square$

### 3.5 Initialization

Theorem 5 and Corollary 6 show that, given a good enough approximation to  $v_1$ , we can rapidly refine this approximation by applying the shifted-and-inverted power method. In this section, we cover initialization. That is, how to obtain a good enough approximation to apply these results.

We first give a simple bound on the quality of a randomly chosen start vector  $x_0$ .

**Lemma 7** (Random Initialization Quality). *Suppose  $x \sim \mathcal{N}(0, \mathbf{I})$ , and we initialize  $x_0$  as  $\frac{x}{\|x\|_2}$ , then with probability greater than  $1 - O\left(\frac{1}{d^{10}}\right)$ , we have:*

$$G(x_0) \leq \sqrt{\kappa(\mathbf{B}^{-1})} d^{10.5} \leq 15 \frac{1}{\sqrt{\text{gap}}} \cdot d^{10.5}$$

where  $\kappa(\mathbf{B}^{-1}) = \lambda_1(\mathbf{B}^{-1})/\lambda_d(\mathbf{B}^{-1})$ .

*Proof.*

$$\begin{aligned}
G(x_0) = G(x) &= \frac{\|\mathbf{P}_{v_1^\perp}(x)\|_{\mathbf{B}}}{\|\mathbf{P}_{v_1}(x)\|_{\mathbf{B}}} = \frac{\sqrt{\|x\|_{\mathbf{B}}^2 - (v_1^\top \mathbf{B}^{1/2} x)^2}}{|v_1^\top \mathbf{B}^{1/2} x|} = \frac{\sqrt{\sum_{i \geq 2} \frac{(v_i^\top x)^2}{\lambda_i(\mathbf{B}^{-1})}}}{\sqrt{\frac{(v_1^\top x)^2}{\lambda_1(\mathbf{B}^{-1})}}}, \\
&\leq \sqrt{\kappa(\mathbf{B}^{-1})} \cdot \frac{\sqrt{\sum_{i \geq 2} (v_i^\top x)^2}}{|v_1^\top x|}
\end{aligned}$$

Since  $\{v_i^\top x\}_i$  are independent standard normal Gaussian variables. By standard concentration arguments, with probability greater than  $1 - e^{-\Omega(d)}$ , we have  $\sqrt{\sum_{i \geq 2} (v_i^\top x)^2} = O(\sqrt{d})$ . Meanwhile,  $v_1^\top x$  is just a one-dimensional standard Gaussian. It is easy to show  $\mathbb{P}(|v_1^\top x| \leq \frac{1}{d^{1/6}}) = O(\frac{1}{d^{1/6}})$ , which finishes the proof.  $\square$

We now show that we can rapidly decrease our initial error to obtain the required  $G(x) \leq \frac{1}{\sqrt{10}}$  bound for Theorem 5.

**Theorem 8** (Approximate Shifted-and-Inverted Power Method – Burn-In). *Suppose we initialize  $x_0$  as in Lemma 7 and suppose we have access to a subroutine  $\text{solve}(\cdot)$  such that*

$$\mathbb{E} [\|\text{solve}(x) - \mathbf{B}^{-1}x\|_{\mathbf{B}}] \leq \frac{1}{3000\kappa(\mathbf{B}^{-1})d^{21}} \cdot \left\| \frac{1}{\lambda - \text{quot}(x)}x - \mathbf{B}^{-1}x \right\|_{\mathbf{B}}$$

where  $\kappa(\mathbf{B}^{-1}) = \lambda_1(\mathbf{B}^{-1})/\lambda_d(\mathbf{B}^{-1})$ . Then the following procedure,

$$x_t = \text{solve}(x_{t-1}) / \|\text{solve}(x_{t-1})\|$$

after  $T = O(\log d + \log \kappa(\mathbf{B}^{-1}))$  iterations satisfies:

$$G(x_T) \leq \frac{1}{\sqrt{10}},$$

with probability greater than  $1 - O(\frac{1}{d^{1/6}})$ .

*Proof.* As before, we first bound the numerator and denominator of  $G(\hat{x})$  more carefully as follows:

$$\begin{aligned}
\text{Numerator: } \quad \|\mathbf{P}_{v_1^\perp}(\hat{x})\|_{\mathbf{B}} &\leq \|\mathbf{P}_{v_1^\perp}(\mathbf{B}^{-1}x)\|_{\mathbf{B}} + \|\mathbf{P}_{v_1^\perp}(\xi)\|_{\mathbf{B}} \leq \|\mathbf{P}_{v_1^\perp}(\mathbf{B}^{-1}x)\|_{\mathbf{B}} + \|\xi\|_{\mathbf{B}} \\
&= \sqrt{\sum_{i \geq 2} (v_i^T \mathbf{B}^{-1/2} x)^2} + \|\xi\|_{\mathbf{B}} = \sqrt{\sum_{i \geq 2} \alpha_i^2 \lambda_i(\mathbf{B}^{-1})} + \|\xi\|_{\mathbf{B}}, \\
\text{Denominator: } \quad \|\mathbf{P}_{v_1}(\hat{x})\|_{\mathbf{B}} &\geq \|\mathbf{P}_{v_1}(\mathbf{B}^{-1}x)\|_{\mathbf{B}} - \|\mathbf{P}_{v_1}(\xi)\|_{\mathbf{B}} \geq \|\mathbf{P}_{v_1}(\mathbf{B}^{-1}x)\|_{\mathbf{B}} - \|\xi\|_{\mathbf{B}} \\
&= |v_1^T \mathbf{B}^{-1/2} x| - \|\xi\|_{\mathbf{B}} = \alpha_1 \sqrt{\lambda_1(\mathbf{B}^{-1})} - \|\xi\|_{\mathbf{B}}
\end{aligned}$$

We now use the above estimates to bound  $G(\hat{x})$ .

$$\begin{aligned}
G(\hat{x}) &\leq \frac{\sqrt{\sum_{i \geq 2} \alpha_i^2 \lambda_i(\mathbf{B}^{-1})} + \|\xi\|_{\mathbf{B}}}{\alpha_1 \sqrt{\lambda_1(\mathbf{B}^{-1})} - \|\xi\|_{\mathbf{B}}} \leq \frac{\lambda_2(\mathbf{B}^{-1}) \sqrt{\sum_{i \geq 2} \frac{\alpha_i^2}{\lambda_i(\mathbf{B}^{-1})}} + \|\xi\|_{\mathbf{B}}}{\lambda_1(\mathbf{B}^{-1}) \sqrt{\frac{\alpha_1^2}{\lambda_1(\mathbf{B}^{-1})}} - \|\xi\|_{\mathbf{B}}} \\
&= G(x) \frac{\lambda_2(\mathbf{B}^{-1}) + \|\xi\|_{\mathbf{B}} / \sqrt{\sum_{i \geq 2} \frac{\alpha_i^2}{\lambda_i(\mathbf{B}^{-1})}}}{\lambda_1(\mathbf{B}^{-1}) - \|\xi\|_{\mathbf{B}} / \sqrt{\frac{\alpha_1^2}{\lambda_1(\mathbf{B}^{-1})}}}
\end{aligned}$$



By Lemma 7, we know with at least probability  $1 - O(\frac{1}{d^{10}})$ , we have  $G(x_0) \leq \sqrt{\kappa(\mathbf{B}^{-1})}d^{10.5}$ .

Conditioned on high probability result of  $G(x_0)$ , we now use induction to prove  $G(x_t) \leq G(x_0)$ . It trivially holds for  $t = 0$ . Suppose we now have  $G(x) \leq G(x_0)$ , then by the condition in Theorem 8 and Markov inequality, we know with probability greater than  $1 - \frac{1}{100\sqrt{\kappa(\mathbf{B}^{-1})}d^{10.5}}$  we have:

$$\begin{aligned} \|\xi\|_B &\leq \frac{1}{30\sqrt{\kappa(\mathbf{B}^{-1})}d^{10.5}} \cdot \left\| \frac{1}{\lambda - \text{quot}(x)}x - \mathbf{B}^{-1}x \right\|_{\mathbf{B}} \\ &\leq \frac{1}{30} \cdot \left\| \frac{1}{\lambda - \text{quot}(x)}x - \mathbf{B}^{-1}x \right\|_{\mathbf{B}} \min \left\{ 1, \frac{1}{G(x_0)} \right\} \\ &\leq \frac{1}{30} \cdot \left\| \frac{1}{\lambda - \text{quot}(x)}x - \mathbf{B}^{-1}x \right\|_{\mathbf{B}} \min \left\{ 1, \frac{1}{G(x)} \right\} \\ &\leq \frac{\lambda_1(\mathbf{B}^{-1}) - \lambda_2(\mathbf{B}^{-1})}{4} \min \left\{ \sqrt{\sum_{i \geq 2} \frac{\alpha_i^2}{\lambda_i(\mathbf{B}^{-1})}}, \sqrt{\frac{\alpha_1^2}{\lambda_1(\mathbf{B}^{-1})}} \right\} \end{aligned}$$

The last inequality uses Corollary 6 with the fact that  $\lambda_2(\mathbf{B}^{-1}) \leq \frac{1}{100}\lambda_1(\mathbf{B}^{-1})$ . Therefore, we have: We will have:

$$G(\hat{x}) \leq \frac{\lambda_1(\mathbf{B}^{-1}) + 3\lambda_2(\mathbf{B}^{-1})}{3\lambda_1(\mathbf{B}^{-1}) + \lambda_2(\mathbf{B}^{-1})} \times G(x) \leq \frac{1}{2}G(x)$$

This finishes the proof of induction.

Finally, by union bound, we know with probability greater than  $1 - O(\frac{1}{d^{10}})$  in  $T = O(\log d + \log \kappa(\mathbf{B}^{-1}))$  steps, we have:

$$G(x_T) \leq \frac{1}{2^T}G(x_0) \leq \frac{1}{\sqrt{10}}$$

□

## 4 Offline Eigenvector Computation

In this section we show how to instantiate the framework of Section 3 in order to compute an approximate top eigenvector in the offline setting. As discussed, in the offline setting we can trivially compute the Rayleigh quotient of a vector in  $\text{nnz}(\mathbf{A})$  time as we have explicit access to  $\mathbf{A}^\top \mathbf{A}$ . Consequently the bulk of our work in this section is to show how we can solve linear systems in  $\mathbf{B}$  efficiently in expectation, allowing us to apply Corollary 6 of Theorem 5.

In Section 4.1 we first show how Stochastic Variance Reduced Gradient (SVRG) [JZ13] can be adapted to solve linear systems of the form  $\mathbf{B}x = b$ . If we wanted, for example, to solve a linear system in a positive definite matrix like  $\mathbf{A}^\top \mathbf{A}$ , we would optimize the objective function  $f(x) = \frac{1}{2}x^\top \mathbf{A}^\top \mathbf{A}x - b^\top x$ . This function can be written as the sum of  $n$  convex components,  $\psi_i(x) = \frac{1}{2}x^\top (a_i a_i^\top) x + \frac{1}{n}b^\top x$ . In each iteration of traditional gradient descent, one computes the full gradient of  $f(x_i)$  and takes a step in that direction. In stochastic gradient methods, at each iteration, a single component is sampled, and the step direction is based only on the gradient of the sampled component. Hence, we avoid a full gradient computation at each iteration, leading to runtime gains.

Unfortunately, while we have access to the rows of  $\mathbf{A}$  and so can solve systems in  $\mathbf{A}^\top \mathbf{A}$ , it is less clear how to solve systems in  $\mathbf{B} = \lambda \mathbf{I} - \mathbf{A}^\top \mathbf{A}$ . To do this, we will split our function into components of the form  $\psi_i(x) = \frac{1}{2}x^\top (w_i \mathbf{I} - a_i a_i^\top) x + \frac{1}{n}b^\top x$  for some set of weights  $w_i$  with  $\sum_{i \in [n]} w_i = \lambda$ .

Importantly,  $w_i \mathbf{I} - a_i a_i^\top$  may not be positive semidefinite. That is, we are minimizing a sum of functions which is convex, but consists of non-convex components. While recent results for minimizing such functions could be applied directly [SS15, CR15] here we show how to obtain stronger results by using a more general form of SVRG and analyzing the specific properties of our function (i.e. the variance).

Our analysis shows that we can make constant factor progress in solving linear systems in  $\mathbf{B}$  in time  $O\left(\text{nnz}(\mathbf{A}) + \frac{d \text{sr}(\mathbf{A})}{\text{gap}^2}\right)$ . If  $\frac{d \text{sr}(\mathbf{A})}{\text{gap}^2} \leq \text{nnz}(\mathbf{A})$  this gives a runtime proportional to the input size – the best we could hope for. If not, we show in Section 4.2 that it is possible to *accelerate* our system solver, achieving runtime  $\tilde{O}\left(\frac{\sqrt{\text{nnz}(\mathbf{A}) d \text{sr}(\mathbf{A})}}{\text{gap}}\right)$ . This result uses the work of [FGKS15, LMH15] on accelerated approximate proximal point algorithms.

With our solvers in place, in Section 4.3 we pull our results together, showing how to use these solvers in the framework of Section 3 to give faster running times for offline eigenvector computation.

## 4.1 SVRG Based Solver

Here we provide a sampling based algorithm for solving linear systems in  $\mathbf{B}$ . In particular we provide an algorithm for solving the more general problem where we are given a strongly convex function that is a sum of possibly non-convex functions that obey smoothness properties. We provide a general result on bounding the progress of an algorithm that solves such a problem by non-uniform sampling in Theorem 9 and then in the remainder of this section we show how to bound the requisite quantities for solving linear systems in  $\mathbf{B}$ .

**Theorem 9** (SVRG for Sums of Non-Convex Functions). *Consider a set of functions,  $\{\psi_1, \psi_2, \dots, \psi_n\}$ , each mapping  $\mathbb{R}^d \rightarrow \mathbb{R}$ . Let  $f(x) = \sum_i \psi_i(x)$  and let  $x^{\text{opt}} \stackrel{\text{def}}{=} \arg \min_{x \in \mathbb{R}^d} f(x)$ . Suppose we have a probability distribution  $p$  on  $[n]$ , and that starting from some initial point  $x_0 \in \mathbb{R}^d$  in each iteration  $k$  we pick  $i_k \in [n]$  independently with probability  $p_{i_k}$  and let*

$$x_{k+1} := x_k - \frac{\eta}{p_i} (\nabla \psi_i(x_k) - \nabla \psi_i(x_0)) + \eta \nabla f(x_0)$$

for some  $\eta$ . If  $f$  is  $\mu$ -strongly convex and if for all  $x \in \mathbb{R}^d$  we have

$$\sum_{i \in [n]} \frac{1}{p_i} \|\nabla \psi_i(x) - \nabla \psi_i(x^{\text{opt}})\|_2^2 \leq 2\bar{S} [f(x) - f(x^{\text{opt}})], \quad (7)$$

where  $\bar{S}$  is a variance parameter, then for all  $m \geq 1$  we have

$$\mathbb{E} \left[ \frac{1}{m} \sum_{k \in [m]} f(x_k) - f(x^{\text{opt}}) \right] \leq \frac{1}{1 - 2\eta\bar{S}} \left[ \frac{1}{\mu\eta m} + 2\eta\bar{S} \right] \cdot [f(x_0) - f(x^{\text{opt}})]$$

Consequently, if we pick  $\eta$  to be a sufficiently small multiple of  $1/\bar{S}$  then when  $m = O(\bar{S}/\mu)$  we can decrease the error by a constant multiplicative factor in expectation.

*Proof.* We first note that  $\mathbb{E}_{i_k}[x_{k+1} - x_k] = \eta \nabla f(x_k)$ . This is, in each iteration, in expectation, we make a step in the direction of the gradient. Using this fact we have:

$$\begin{aligned} \mathbb{E}_{i_k} \|x_{k+1} - x^{\text{opt}}\|_2^2 &= \mathbb{E}_{i_k} \|(x_{k+1} - x_k) + (x_k - x^{\text{opt}})\|_2^2 \\ &= \|x_k - x^{\text{opt}}\|_2^2 - 2\mathbb{E}_{i_k} (x_{k+1} - x_k)^\top (x_k - x^{\text{opt}}) + \mathbb{E}_{i_k} \|x_{k+1} - x_k\|_2^2 \\ &= \|x_k - x^{\text{opt}}\|_2^2 - 2\eta \nabla f(x_k)^\top (x_k - x^{\text{opt}}) \\ &\quad + \sum_{i \in [n]} \eta^2 p_i \left\| \frac{1}{p_i} (\nabla \psi_i(x_k) - \nabla \psi_i(x_0)) + \nabla f(x_0) \right\|_2^2 \end{aligned}$$

We now apply the fact that  $\|x + y\|_2^2 \leq 2\|x\|_2^2 + 2\|y\|_2^2$  to give:

$$\begin{aligned} &\sum_{i \in [n]} p_i \left\| \frac{1}{p_i} (\nabla \psi_i(x_k) - \nabla \psi_i(x_0)) + \nabla f(x_0) \right\|_2^2 \\ &\leq \sum_{i \in [n]} 2p_i \left\| \frac{1}{p_i} (\nabla \psi_i(x_k) - \nabla \psi_i(x^{\text{opt}})) \right\|_2^2 + \sum_{i \in [n]} 2p_i \left\| \frac{1}{p_i} (\nabla \psi_i(x_0) - \nabla \psi_i(x^{\text{opt}})) - \nabla f(x_0) \right\|_2^2. \end{aligned}$$

Then, using that  $\nabla f(x^{\text{opt}}) = 0$  by optimality, that  $\mathbb{E} \|x - \mathbb{E}x\|_2^2 \leq \mathbb{E} \|x\|_2^2$ , and (7) we have:

$$\begin{aligned} &\sum_{i \in [n]} p_i \left\| \frac{1}{p_i} (\nabla \psi_i(x_k) - \nabla \psi_i(x_0)) + \nabla f(x_0) \right\|_2^2 \\ &\leq \sum_{i \in [n]} \frac{2}{p_i} \|\nabla \psi_i(x_k) - \nabla \psi_i(x^{\text{opt}})\|_2^2 + \sum_{i \in [n]} 2p_i \left\| \frac{1}{p_i} (\nabla \psi_i(x_0) - \nabla \psi_i(x^{\text{opt}})) - (\nabla f(x_0) - \nabla f(x^{\text{opt}})) \right\|_2^2 \\ &\leq \sum_{i \in [n]} \frac{2}{p_i} \|\nabla \psi_i(x_k) - \nabla \psi_i(x^{\text{opt}})\|_2^2 + \sum_{i \in [n]} 2p_i \left\| \frac{1}{p_i} \nabla \psi_i(x_0) - \nabla \psi_i(x^{\text{opt}}) \right\|_2^2 \\ &\leq 4\bar{S} [f(x_k) - f(x^{\text{opt}}) + f(x_0) - f(x^{\text{opt}})] \end{aligned}$$

Since  $f(x^{\text{opt}}) - f(x_k) \geq \nabla f(x_k)^\top (x^{\text{opt}} - x_k)$  by the convexity of  $f$ , these inequalities imply

$$\begin{aligned} \mathbb{E}_{i_k} \|x_{k+1} - x^{\text{opt}}\|_2^2 &\leq \|x_k - x^{\text{opt}}\|_2^2 - 2\eta [f(x_k) - f(x^{\text{opt}})] + 4\eta^2 \bar{S} [f(x_k) - f(x^{\text{opt}}) + f(x_0) - f(x^{\text{opt}})] \\ &= \|x_k - x^{\text{opt}}\|_2^2 - 2\eta(1 - 2\eta\bar{S}) (f(x_k) - f(x^{\text{opt}})) + 4\eta^2 \bar{S} (f(x_0) - f(x^{\text{opt}})) \end{aligned}$$

Rearranging, we have:

$$2\eta(1 - 2\eta\bar{S}) (f(x_k) - f(x^{\text{opt}})) \leq \|x_k - x^{\text{opt}}\|_2^2 - \mathbb{E}_{i_k} \|x_{k+1} - x^{\text{opt}}\|_2^2 + 4\eta^2 \bar{S} (f(x_0) - f(x^{\text{opt}})).$$

And summing over all iterations and taking expectations we have:

$$\mathbb{E} \left[ 2\eta(1 - 2\eta\bar{S}) \sum_{k \in [m]} f(x_k) - f(x^{\text{opt}}) \right] \leq \|x_0 - x^{\text{opt}}\|_2^2 + 4m\eta^2 \bar{S} [f(x_0) - f(x^{\text{opt}})].$$

Finally, we use that by strong convexity,  $\|x_0 - x^{\text{opt}}\|_2^2 \leq \frac{2}{\mu} (f(x_0) - f(x^{\text{opt}}))$  to obtain:

$$\mathbb{E} \left[ 2\eta(1 - 2\eta\bar{S}) \sum_{k \in [m]} f(x_k) - f(x^{\text{opt}}) \right] \leq \frac{2}{\mu} [f(x_0) - f(x^{\text{opt}})] + 4m\eta^2\bar{S} [f(x_0) - f(x^{\text{opt}})]$$

and thus

$$\mathbb{E} \left[ \frac{1}{m} \sum_{k \in [m]} f(x_k) - f(x^{\text{opt}}) \right] \leq \frac{1}{1 - 2\eta\bar{S}} \left[ \frac{1}{\mu\eta m} + 2\eta\bar{S} \right] \cdot [f(x_0) - f(x^{\text{opt}})]$$

□

Theorem 9 immediately yields a solver for  $\mathbf{B}x = b$ . Finding the minimum norm solution to this system is equivalent to minimizing  $f(x) = \frac{1}{2}x^\top \mathbf{B}x - b^\top x$ . If we take the common approach of applying a smoothness bound for each  $\psi_i$  along with a strong convexity bound on  $f(x)$  we obtain:

**Lemma 10** (Simple Variance Bound for SVRG). *Let*

$$\psi_i(x) \stackrel{\text{def}}{=} \frac{1}{2}x^\top \left( \frac{\lambda \|a_i\|_2^2}{\|\mathbf{A}\|_F^2} \mathbf{I} - a_i a_i^\top \right) x - \frac{1}{n} b^\top x$$

so we have  $\sum_{i \in [n]} \psi_i(x) = f(x) = \frac{1}{2}x^\top \mathbf{B}x - b^\top x$ . Setting  $p_i = \frac{\|a_i\|_2^2}{\|\mathbf{A}\|_F^2}$  for all  $i$ , we have

$$\sum_{i \in [n]} \frac{1}{p_i} \|\nabla \psi_i(x) - \nabla \psi_i(x^{\text{opt}})\|_2^2 = O \left( \frac{\|\mathbf{A}\|_F^4}{\lambda - \lambda_1} [f(x) - f(x^{\text{opt}})] \right)$$

*Proof.* We first compute, for all  $i \in [n]$

$$\nabla \psi_i(x) = \left( \frac{\lambda \|a_i\|_2^2}{\|\mathbf{A}\|_F^2} \mathbf{I} - a_i a_i^\top \right) x - \frac{1}{n} b. \quad (8)$$

We have that each  $\psi_i$  is  $\frac{\lambda \|a_i\|_2^2}{\|\mathbf{A}\|_F^2} + \|a_i\|^2$  smooth with respect to  $\|\cdot\|_2$ . Specifically,

$$\begin{aligned} \|\nabla \psi_i(x) - \nabla \psi_i(x^{\text{opt}})\|_2 &= \left\| \left( \frac{\lambda \|a_i\|_2^2}{\|\mathbf{A}\|_F^2} \mathbf{I} - a_i a_i^\top \right) (x - x^{\text{opt}}) \right\|_2 \\ &\leq \left( \frac{\lambda \|a_i\|_2^2}{\|\mathbf{A}\|_F^2} + \|a_i\|^2 \right) \|x - x^{\text{opt}}\|_2. \end{aligned}$$

Additionally,  $f(x)$  is  $\lambda_d(\mathbf{B}) = \lambda - \lambda_1$  strongly convex so we have  $\|x - x^{\text{opt}}\|_2^2 \leq \frac{2}{\lambda - \lambda_1} [f(x) - f(x^{\text{opt}})]$  and putting all this together we have

$$\begin{aligned} \sum_{i \in [n]} \frac{1}{p_i} \|\nabla \psi_i(x) - \nabla \psi_i(x^{\text{opt}})\|_2^2 &\leq \sum_{i \in [n]} \frac{\|\mathbf{A}\|_F^2}{\|a_i\|_2^2} \cdot \|a_i\|_2^4 \left( \frac{\lambda}{\|\mathbf{A}\|_F^2} + 1 \right)^2 \cdot \frac{2}{\lambda - \lambda_1} [f(x) - f(x^{\text{opt}})] \\ &= O \left( \frac{\|\mathbf{A}\|_F^4}{\lambda - \lambda_1} [f(x) - f(x^{\text{opt}})] \right) \end{aligned}$$

where the last step uses that  $\lambda \leq 2\lambda_1 \leq 2\|\mathbf{A}\|_F^2$  so  $\frac{\lambda}{\|\mathbf{A}\|_F^2} \leq 2$ .  $\square$

Assuming that  $\lambda = (1 + c \cdot \text{gap})\lambda_1$  for some constant  $c$ , the above bound means that we can make constant progress on our linear system by setting  $m = O(\bar{S}/\mu) = O\left(\frac{\|\mathbf{A}\|_F^4}{(\lambda - \lambda_1)^2}\right) = O\left(\frac{\text{sr}(\mathbf{A})^2}{\text{gap}^2}\right)$ . This dependence on stable rank matches the dependence given in [Sha15c] (see discussion in Section 1.3), however we can show that it is suboptimal. We show to improve the bound to  $O\left(\frac{\text{sr}(\mathbf{A})}{\text{gap}^2}\right)$  by using a better variance analysis. Instead of bounding each  $\|\nabla\psi_i(x) - \nabla\psi_i(x^{\text{opt}})\|_2^2$  term using the smoothness of  $\psi_i$ , we more carefully bound the sum of these terms.

**Lemma 11.** (*Improved Variance Bound for SVRG*) For  $i \in [n]$  let

$$\psi_i(x) \stackrel{\text{def}}{=} \frac{1}{2}x^\top \left( \frac{\lambda \|a_i\|_2^2}{\|\mathbf{A}\|_F^2} \mathbf{I} - a_i a_i^\top \right) x - \frac{1}{n} b^\top x$$

so we have  $\sum_{i \in [n]} \psi_i(x) = f(x) = \frac{1}{2}x^\top \mathbf{B}x - b^\top x$ . Setting  $p_i = \frac{\|a_i\|_2^2}{\|\mathbf{A}\|_F^2}$  for all  $i$ , we have for all  $x$

$$\sum_{i \in [n]} \frac{1}{p_i} \|\nabla\psi_i(x) - \nabla\psi_i(x^{\text{opt}})\|_2^2 \leq \frac{4\lambda_1 \|\mathbf{A}\|_F^2}{\lambda - \lambda_1} \cdot [f(x) - f(x^{\text{opt}})].$$

*Proof.* Using the gradient computation in (8) we have

$$\begin{aligned} \sum_{i \in [n]} \frac{1}{p_i} \|\nabla\psi_i(x) - \nabla\psi_i(x^{\text{opt}})\|_2^2 &= \sum_{i \in [n]} \frac{\|\mathbf{A}\|_F^2}{\|a_i\|_2^2} \left\| \left( \frac{\lambda \|a_i\|_2^2}{\|\mathbf{A}\|_F^2} \mathbf{I} - a_i a_i^\top \right) (x - x^{\text{opt}}) \right\|_2^2 \\ &= \sum_{i \in [n]} \frac{\lambda^2 \|a_i\|_2^2}{\|\mathbf{A}\|_F^2} \|x - x^{\text{opt}}\|_2^2 - 2 \sum_{i \in [n]} \lambda \|x - x^{\text{opt}}\|_{a_i a_i^\top}^2 \\ &\quad + \sum_{i \in [n]} \frac{\|\mathbf{A}\|_F^2}{\|a_i\|_2^2} \|x - x^{\text{opt}}\|_{\|a_i\|_2^2 a_i a_i^\top}^2 \\ &= \lambda^2 \|x - x^{\text{opt}}\|_2^2 - 2\lambda \|x - x^{\text{opt}}\|_\Sigma^2 + \|\mathbf{A}\|_F^2 \|x - x^{\text{opt}}\|_\Sigma^2 \\ &\leq \lambda \|x - x^{\text{opt}}\|_\mathbf{B}^2 + \|\mathbf{A}\|_F^2 \|x - x^{\text{opt}}\|_\Sigma^2. \end{aligned} \tag{9}$$

Now since

$$\Sigma \preceq \lambda_1 \mathbf{I} \preceq \frac{\lambda_1}{\lambda - \lambda_1} \mathbf{B}$$

we have

$$\begin{aligned} \sum_{i \in [n]} \frac{1}{p_i} \|\nabla\psi_i(x) - \nabla\psi_i(x^{\text{opt}})\|_2^2 &\leq \left( \frac{\lambda(\lambda - \lambda_1) + \|\mathbf{A}\|_F^2 \cdot \lambda_1}{\lambda - \lambda_1} \right) \|x - x^{\text{opt}}\|_\mathbf{B}^2 \\ &\leq \left( \frac{2\|\mathbf{A}\|_F^2 \lambda_1}{\lambda - \lambda_1} \right) \|x - x^{\text{opt}}\|_\mathbf{B}^2 \end{aligned}$$

where in the last inequality we just coarsely bound  $\lambda(\lambda - \lambda_1) \leq \lambda_1 \|\mathbf{A}\|_F^2$ . Now since  $\mathbf{B}$  is full rank,  $\mathbf{B}x^{\text{opt}} = b$ , we can compute:

$$\|x - x^{\text{opt}}\|_{\mathbf{B}}^2 = x^\top \mathbf{B}x - 2b^\top x + b^\top x^{\text{opt}} = 2[f(x) - f(x^{\text{opt}})]. \quad (10)$$

The result follows.  $\square$

Plugging the bound in Lemma 11 into Theorem 9 we have:

**Theorem 12.** (*Offline SVRG-Based Solver*) Let  $\bar{S} = \frac{2\lambda_1 \|\Sigma\|_F^2}{\lambda - \lambda_1}$ ,  $\mu = \lambda - \lambda_1$ . The iterative procedure described in Theorem 9 with  $f(x) = \frac{1}{2}x^\top \mathbf{B}x - b^\top x$ ,  $\psi_i(x) = \frac{1}{2}x^\top \left( \frac{\lambda \|a_i\|_2^2}{\|\Sigma\|_F^2} \mathbf{I} - a_i a_i^\top \right) x - b^\top x$ ,  $p_i = \frac{\|a_i\|_2^2}{\|\Sigma\|_F^2}$ ,  $\eta = 1/(8\bar{S})$  and  $m$  chosen uniformly at random from  $[64\bar{S}/\mu]$  returns a vector  $x_m$  such that

$$\mathbb{E} \|x_m - x^{\text{opt}}\|_{\mathbf{B}}^2 \leq \frac{1}{2} \|x_0 - x^{\text{opt}}\|_{\mathbf{B}}^2.$$

Further, assuming  $(1 + \frac{\text{gap}}{150}) \lambda_1 < \lambda \leq (1 + \frac{\text{gap}}{100}) \lambda_1$ , this procedure runs in time  $O\left(\text{nnz}(\mathbf{A}) + \frac{d \cdot \text{sr}(\mathbf{A})}{\text{gap}^2}\right)$ .

*Proof.* Lemma 11 tells us that

$$\sum_{i \in [n]} \frac{1}{p_i} \|\nabla \psi_i(x) - \nabla \psi_i(x^{\text{opt}})\|_2^2 \leq 2\bar{S} [f(x) - f(x^{\text{opt}})].$$

Further  $f(x) = \frac{1}{2}x^\top \mathbf{B}x - b^\top x$  is  $\lambda_d(\mathbf{B})$ -strongly convex and  $\lambda_d(\mathbf{B}) = \lambda - \lambda_1 = \mu$ . Plugging this into Theorem 9 and using (10) which shows  $\|x - x^{\text{opt}}\|_{\mathbf{B}}^2 = 2[f(x) - f(x^{\text{opt}})]$  we have, for  $m$  chosen uniformly from  $[64\bar{S}/\mu]$ :

$$\begin{aligned} \mathbb{E} \left[ \frac{1}{64\bar{S}/\mu} \sum_{k \in [64\bar{S}/\mu]} f(x_k) - f(x^{\text{opt}}) \right] &\leq 4/3 \cdot [1/8 + 1/8] \cdot [f(x_0) - f(x^{\text{opt}})] \\ \mathbb{E} [f(x_m) - f(x^{\text{opt}})] &\leq \frac{1}{2} [f(x_0) - f(x^{\text{opt}})] \\ \mathbb{E} \|x_m - x^{\text{opt}}\|_{\mathbf{B}}^2 &\leq \frac{1}{2} \|x_0 - x^{\text{opt}}\|_{\mathbf{B}}^2. \end{aligned}$$

The procedure requires  $O(\text{nnz}(\mathbf{A}))$  time to initially compute  $\nabla f(x_0)$ , along with each  $p_i$  and the step size  $\eta$  which depend on  $\|\mathbf{A}\|_F^2$  and the row norms of  $\mathbf{A}$ . Each iteration then just requires  $O(d)$  time to compute  $\nabla \psi_i(\cdot)$  and perform the necessary vector operations. Since there are at most  $[64\bar{S}/\mu] = O\left(\frac{\lambda_1 \|\mathbf{A}\|_F^2}{(\lambda - \lambda_1)^2}\right)$  iterations, our total runtime is

$$O\left(\text{nnz}(\mathbf{A}) + d \cdot \frac{\lambda_1 \|\mathbf{A}\|_F^2}{(\lambda - \lambda_1)^2}\right) = O\left(\text{nnz}(\mathbf{A}) + \frac{d \cdot \text{sr}(\mathbf{A})}{\text{gap}^2}\right).$$

Note that if our matrix is uniformly sparse - i.e. all rows have sparsity at most  $d_s$ , then the runtime is actually at most  $O\left(\text{nnz}(\mathbf{A}) + \frac{d_s \cdot \text{sr}(\mathbf{A})}{\text{gap}^2}\right)$ .  $\square$

## 4.2 Accelerated Solver

Theorem 12 gives a linear solver for  $\mathbf{B}$  that makes progress in expectation and which we can plug into Theorems 5 and 8. However, we first show that the runtime in Theorem 12 can be accelerated in some cases. We apply a result of [FGKS15], which shows that, given a solver for a regularized version of a convex function  $f(x)$ , we can produce a fast solver for  $f(x)$  itself. Specifically:

**Lemma 13** (Theorem 1.1 of [FGKS15]). *Let  $f(x)$  be a  $\mu$ -strongly convex function and let  $x^{\text{opt}} \stackrel{\text{def}}{=} \arg \min_{x \in \mathbb{R}^d} f(x)$ . For any  $\gamma > 0$  and any  $x_0 \in \mathbb{R}^d$ , let  $f_{\gamma, x_0}(x) \stackrel{\text{def}}{=} f(x) + \frac{\gamma}{2} \|x - x_0\|_2^2$ . Let  $x_{\gamma, x_0}^{\text{opt}} \stackrel{\text{def}}{=} \arg \min_{x \in \mathbb{R}^d} f_{\gamma, x_0}(x)$ .*

*Suppose that, for all  $x_0 \in \mathbb{R}^d$ ,  $c > 0$ ,  $\gamma > 0$ , we can compute a point  $x_c$  such that*

$$\mathbb{E} f_{\gamma, x_0}(x_c) - f_{\gamma, x_0}(x_{\gamma, x_0}^{\text{opt}}) \leq \frac{1}{c} [f_{\gamma, x_0} - f_{\gamma, x_0}(x_{\gamma, x_0}^{\text{opt}})]$$

*in time  $\mathcal{T}_c$ . Then given any  $x_0$ ,  $c > 0$ ,  $\gamma > 2\mu$ , we can compute  $x_1$  such that*

$$\mathbb{E} f(x_1) - f(x^{\text{opt}}) \leq \frac{1}{c} [f(x_0) - f(x^{\text{opt}})]$$

*in time  $O\left(\mathcal{T}_{4\left(\frac{2\gamma+\mu}{\mu}\right)^{3/2} \sqrt{\lceil \gamma/\mu \rceil} \log c}\right)$ .*

We first give a new variance bound on solving systems in  $\mathbf{B}$  when a regularizer is used. The proof of this bound is very close to the proof given for the unregularized problem in Lemma 11.

**Lemma 14.** *For  $i \in [n]$  let*

$$\psi_i(x) \stackrel{\text{def}}{=} \frac{1}{2} x^\top \left( \frac{\lambda \|a_i\|_2^2}{\|\mathbf{A}\|_F^2} \mathbf{I} - a_i a_i^\top \right) x - \frac{1}{n} b^\top x + \frac{\gamma \|a_i\|_2^2}{2 \|\mathbf{A}\|_F^2} \|x - x_0\|_2^2$$

*so we have  $\sum_{i \in [n]} \psi_i(x) = f_{\gamma, x_0}(x) = \frac{1}{2} x^\top \mathbf{B} x - b^\top x + \frac{\gamma}{2} \|x - x_0\|_2^2$ . Setting  $p_i = \frac{\|a_i\|_2^2}{\|\mathbf{A}\|_F^2}$  for all  $i$ , we have for all  $x$*

$$\sum_{i \in [n]} \frac{1}{p_i} \|\nabla \psi_i(x) - \nabla \psi_i(x_{\gamma, x_0}^{\text{opt}})\|_2^2 \leq \left( \frac{\gamma^2 + 12\lambda_1 \|\mathbf{A}\|_F^2}{\lambda - \lambda_1 + \gamma} \right) [f_{\gamma, x_0}(x) - f_{\gamma, x_0}(x_{\gamma, x_0}^{\text{opt}})]$$

*Proof.* We have for all  $i \in [n]$

$$\nabla \psi_i(x) = \left( \frac{\lambda \|a_i\|_2^2}{\|\mathbf{A}\|_F^2} \mathbf{I} - a_i a_i^\top \right) x - \frac{1}{n} b + \frac{\gamma \|a_i\|_2^2}{2 \|\mathbf{A}\|_F^2} (x - 2x_0) \quad (11)$$

Plugging this in we have:

$$\sum_{i \in [n]} \frac{1}{p_i} \|\nabla \psi_i(x) - \nabla \psi_i(x_{\gamma, x_0}^{\text{opt}})\|_2^2 = \sum_{i \in [n]} \frac{\|\mathbf{A}\|_F^2}{\|a_i\|_2^2} \left\| \left( \frac{\lambda \|a_i\|_2^2}{\|\mathbf{A}\|_F^2} \mathbf{I} - a_i a_i^\top \right) (x - x_{\gamma, x_0}^{\text{opt}}) + \frac{\gamma \|a_i\|_2^2}{2 \|\mathbf{A}\|_F^2} (x - x_{\gamma, x_0}^{\text{opt}}) \right\|_2^2$$

For simplicity we now just use the fact that  $\|x + y\|_2^2 \leq 2\|x\|_2^2 + 2\|y\|_2^2$  and apply our bound from equation (9) to obtain:

$$\begin{aligned}
\sum_{i \in [n]} \frac{1}{p_i} \|\nabla \psi_i(x) - \nabla \psi_i(x_{\gamma, x_0}^{\text{opt}})\|_2^2 &\leq 2\lambda^2 \|x - x_{\gamma, x_0}^{\text{opt}}\|_2^2 - 4\lambda \|x - x_{\gamma, x_0}^{\text{opt}}\|_{\Sigma}^2 + 2\|\Sigma\|_F^2 \|x - x_{\gamma, x_0}^{\text{opt}}\|_{\Sigma}^2 \\
&\quad + 2 \sum_{i \in [n]} \frac{\|a_i\|_2^2 \gamma^2}{\|\mathbf{A}\|_F^2 4} \|x - x_{\gamma, x_0}^{\text{opt}}\|_2^2 \\
&\leq \left(2\lambda^2 + \gamma^2/2 + 2\lambda_1 \|\mathbf{A}\|_F^2 - 4\lambda_1 \lambda\right) \|x - x_{\gamma, x_0}^{\text{opt}}\|_2^2 \\
&\leq \left(\gamma^2/2 + 6\lambda_1 \|\mathbf{A}\|_F^2\right) \|x - x_{\gamma, x_0}^{\text{opt}}\|_2^2
\end{aligned}$$

Now,  $f_{\gamma, x_0}(\cdot)$  is  $\lambda - \lambda_1 + \gamma$  strongly convex, so

$$\|x - x_{\gamma, x_0}^{\text{opt}}\|_2^2 \leq \frac{2}{\lambda - \lambda_1 + \gamma} [f_{\gamma, x_0}(x) - f_{\gamma, x_0}(x_{\gamma, x_0}^{\text{opt}})].$$

So overall we have:

$$\sum_{i \in [n]} \frac{1}{p_i} \|\nabla \psi_i(x) - \nabla \psi_i(x_{\gamma, x_0}^{\text{opt}})\|_2^2 \leq \left( \frac{\gamma^2 + 12\lambda_1 \|\mathbf{A}\|_F^2}{\lambda - \lambda_1 + \gamma} \right) [f_{\gamma, x_0}(x) - f_{\gamma, x_0}(x_{\gamma, x_0}^{\text{opt}})]$$

□

We can now use this variance bound to obtain an accelerated solver for  $\mathbf{B}$ . We assume  $\text{nnz}(\mathbf{A}) \leq \frac{d_{\text{sr}}(\mathbf{A})}{\text{gap}^2}$ , as otherwise, the unaccelerated solver in Theorem 12 runs in  $O(\text{nnz}(\mathbf{A}))$  time and cannot be accelerated further.

**Theorem 15** (Accelerated SVRG-Based Solver). *Assuming  $(1 + \frac{\text{gap}}{150}) \lambda_1 < \lambda \leq (1 + \frac{\text{gap}}{100}) \lambda_1$  and  $\text{nnz}(\mathbf{A}) \leq \frac{d_{\text{sr}}(\mathbf{A})}{\text{gap}^2}$ , applying the iterative procedure described in Theorem 9 along with the acceleration given by Lemma 13 gives a solver that returns  $x$  with*

$$\mathbb{E} \|x - x^{\text{opt}}\|_{\mathbf{B}}^2 \leq \frac{1}{2} \|x_0 - x^{\text{opt}}\|_{\mathbf{B}}^2.$$

in time  $O\left(\frac{\text{nnz}(\mathbf{A})^{3/4} (d_{\text{sr}}(\mathbf{A}))^{1/4}}{\sqrt{\text{gap}}} \cdot \log\left(\frac{d}{\text{gap}}\right)\right)$ .

*Proof.* Following Theorem 12, the variance bound of Lemma 14 means that we can make constant progress in minimizing  $f_{\gamma, x_0}(x)$  in  $O(\text{nnz}(\mathbf{A}) + dm)$  time where  $m = O\left(\frac{\gamma^2 + 12\lambda_1 \|\Sigma\|_F^2}{(\lambda - \lambda_1 + \gamma)^2}\right)$ . So, for  $\gamma \geq 2(\lambda - \lambda_1)$  we can make  $4\left(\frac{2\gamma + (\lambda - \lambda_1)}{\lambda - \lambda_1}\right)^{3/2}$  progress, as required by Lemma 13 in time  $O\left((\text{nnz}(\mathbf{A}) + dm) \cdot \log\left(\frac{\gamma}{\lambda - \lambda_1}\right)\right)$  time. Hence by Lemma 13 we can make constant factor expected progress in minimizing  $f(x)$  in time:

$$O\left(\left(\text{nnz}(\mathbf{A}) + d \frac{\gamma^2 + 12\lambda_1 \|\mathbf{A}\|_F^2}{(\lambda - \lambda_1 + \gamma)^2}\right) \log\left(\frac{\gamma}{\lambda - \lambda_1}\right) \sqrt{\frac{\gamma}{\lambda - \lambda_1}}\right)$$



By our assumption, we have  $\text{nnz}(\mathbf{A}) \leq \frac{d \text{sr}(\mathbf{A})}{\text{gap}^2} = \frac{d\lambda_1 \|\mathbf{A}\|_F^2}{(\lambda - \lambda_1)^2}$ . So, if we let  $\gamma = \Theta\left(\sqrt{\frac{d\lambda_1 \|\mathbf{A}\|_F^2}{\text{nnz}(\mathbf{A})}}\right)$  then using a sufficiently large constant, we have  $\gamma \geq 2(\lambda - \lambda_1)$ . We have  $\frac{\gamma}{\lambda - \lambda_1} = \Theta\left(\sqrt{\frac{d\lambda_1 \|\mathbf{A}\|_F^2}{\text{nnz}(\mathbf{A})\lambda_1^2 \text{gap}^2}}\right) = \Theta\left(\sqrt{\frac{d \text{sr}(\mathbf{A})}{\text{nnz}(\mathbf{A})\text{gap}^2}}\right)$  and can make constant expected progress in minimizing  $f(x)$  in time:

$$O\left(\frac{\text{nnz}(\mathbf{A})^{3/4} (d \text{sr}(\mathbf{A}))^{1/4}}{\sqrt{\text{gap}}} \cdot \log\left(\frac{d}{\text{gap}}\right)\right).$$

□

### 4.3 Shifted-and-Inverted Power Method

Finally, we are able to combine the solvers from Sections 4.1 and 4.2 with the framework of Section 3 to obtain faster algorithms for top eigenvector computation.

**Theorem 16** (Shifted-and-Inverted Power Method With SVRG). *Let  $\mathbf{B} = \lambda\mathbf{I} - \mathbf{A}^\top \mathbf{A}$  for  $(1 + \frac{\text{gap}}{150})\lambda_1 \leq \lambda \leq (1 + \frac{\text{gap}}{100})\lambda_1$  and let  $x_0 \sim \mathcal{N}(0, \mathbf{I})$  be a random initial vector. Running the inverted power method on  $\mathbf{B}$  initialized with  $x_0$ , using the SVRG solver from Theorem 12 to approximately apply  $\mathbf{B}^{-1}$  at each step, returns  $x$  such that with probability  $1 - O(\frac{1}{d^{10}})$ ,  $x^\top \Sigma x \geq (1 - \epsilon)\lambda_1$  in total time*

$$O\left(\left(\text{nnz}(\mathbf{A}) + \frac{d \text{sr}(\mathbf{A})}{\text{gap}^2}\right) \cdot \left(\log^2\left(\frac{d}{\text{gap}}\right) + \log\left(\frac{1}{\epsilon}\right)\right)\right).$$

Note that by instantiating the above theorem with  $\epsilon' = \epsilon \cdot \text{gap}$ , and applying Lemma 1 we can find  $x$  such that  $|v_1^\top x| \geq 1 - \epsilon$  in the same asymptotic running time (an extra  $\log(1/\text{gap})$  term is absorbed into the  $\log^2(d/\text{gap})$  term).

*Proof.* By Theorem 8, if we start with  $x_0 \sim \mathcal{N}(0, \mathbf{I})$  we can run  $O\left(\log\left(\frac{d}{\text{gap}}\right)\right)$  iterations of the inverted power method, to obtain  $x_1$  with  $G(x_1) \leq \frac{1}{\sqrt{10}}$  with probability  $1 - O(\frac{1}{d^{10}})$ . Each iteration requires applying an linear solver that decreases initial error in expectation by a factor of  $\frac{1}{\text{poly}(d, 1/\text{gap})}$ . Such a solver is given by applying the solver in Theorem 12  $O\left(\log\left(\frac{d}{\text{gap}}\right)\right)$  times, decreasing error by a constant factor in expectation each time. So overall in order to find  $x_1$  with  $G(x_1) \leq \frac{1}{\sqrt{10}}$ , we require time  $O\left(\left(\text{nnz}(\mathbf{A}) + \frac{d \text{sr}(\mathbf{A})}{\text{gap}^2}\right) \cdot \log^2\left(\frac{d}{\text{gap}}\right)\right)$ .

After this initial ‘burn-in’ period we can apply Corollary 6 of Theorem 5, which shows that running a single iteration of the inverted power method will decrease  $G(x)$  by a constant factor in expectation. In such an iteration, we only need to use a solver that decreases initial error by a constant factor in expectation. So we can perform each inverted power iteration in this stage in time  $O\left(\text{nnz}(\mathbf{A}) + \frac{d \text{sr}(\mathbf{A})}{\text{gap}^2}\right)$ .

With  $O\left(\log\left(\frac{d}{\epsilon}\right)\right)$  iterations, we can obtain  $x$  with  $\mathbb{E}[G(x)^2] = O\left(\frac{\epsilon}{d^{10}}\right)$ . So by Markov’s inequality, we have  $G(x)^2 = O(\epsilon)$ , giving us  $x^\top \Sigma x \geq (1 - O(\epsilon))\lambda_1$  by Lemma 3. Union bounding over both stages gives us failure probability  $O\left(\frac{1}{d^{10}}\right)$ , and adding the runtimes from the two stages gives us the final result. Note that the second stage requires  $O\left(\log\left(\frac{d}{\epsilon}\right)\right) = O(\log d + \log(1/\epsilon))$  iterations to achieve the high probability bound. However, the  $O(\log d)$  term is smaller than the  $O\left(\log^2\left(\frac{d}{\text{gap}}\right)\right)$  term, so is absorbed into the asymptotic notation.

□

We can apply an identical analysis using the accelerated solver from Theorem 15, obtaining the following runtime which beats Theorem 16 whenever  $\text{nnz}(\mathbf{A}) \leq \frac{d \text{sr}(\mathbf{A})}{\text{gap}^2}$ :

**Theorem 17** (Shifted-and-Inverted Power Method Using Accelerated SVRG). *Let  $\mathbf{B} = \lambda \mathbf{I} - \mathbf{A}^\top \mathbf{A}$  for  $(1 + \frac{\text{gap}}{150}) \lambda_1 \leq \lambda \leq (1 + \frac{\text{gap}}{100}) \lambda_1$  and let  $x_0 \sim \mathcal{N}(0, \mathbf{I})$  be a random initial vector. Assume that  $\text{nnz}(\mathbf{A}) \leq \frac{d \text{sr}(\mathbf{A})}{\text{gap}^2}$ . Running the inverted power method on  $\mathbf{B}$  initialized with  $x_0$ , using the accelerated SVRG solver from Theorem 15 to approximately apply  $\mathbf{B}^{-1}$  at each step, returns  $x$  such that with probability  $1 - O(\frac{1}{d^{10}})$ ,  $|v_1^\top x| \geq 1 - \epsilon$  in total time*

$$O\left(\left(\frac{\text{nnz}(\mathbf{A})^{3/4}(d \text{sr}(\mathbf{A}))^{1/4}}{\sqrt{\text{gap}}}\right) \cdot \left(\log^3\left(\frac{d}{\text{gap}}\right) + \log\left(\frac{d}{\text{gap}}\right) \log\left(\frac{1}{\epsilon}\right)\right)\right).$$

## 5 Statistical Setting

Here we show how to apply the shifted-and-inverted power method framework of Section 3 to the online setting. This setting is more difficult than the offline case. As there is no canonical matrix  $\mathbf{A}$ , and we only have access to the distribution  $\mathcal{D}$  through samples, in order to apply Theorem 5 we must show how to both estimate the Rayleigh quotient (Section 5.1) as well as solve the requisite linear systems in expectation (Section 5.2).

After laying this ground work, our main result is given in Section 5.3. Ultimately, the results in this section allow us to achieve more efficient algorithms for computing the top eigenvector in the statistical setting as well as improve upon the previous best known sample complexity for top eigenvector computation. As we show in Section 7 the bounds we provide in this section are in fact tight for general distributions.

### 5.1 Estimating the Rayleigh Quotient

Here we show how to estimate the Rayleigh quotient of a vector with respect to  $\Sigma$ . Our analysis is standard – we first approximate the Rayleigh quotient by its empirical value on a batch of  $k$  samples and prove using Chebyshev’s inequality that the error on this sample is small with constant probability. We then repeat this procedure  $O(\log(1/p))$  times and output the median. By Chernoff bound this yields a good estimate with probability  $1 - p$ . The formal statement of this result and its proof comprise the remainder of this subsection.

**Theorem 18** (Online Rayleigh Quotient Estimation). *Given any  $\epsilon \in (0, 1]$  and  $p \in [0, 1]$  let  $k = \lceil 4 v(\mathcal{D}) \epsilon^{-2} \rceil$  and  $m = c \log(1/p)$  for some sufficiently large constant  $c$ . For all  $i \in [k]$  and  $j \in [m]$  let  $a_i^{(j)}$  be drawn independently from  $\mathcal{D}$ . Then if for any unit norm  $x$  we let  $R_{i,j} \stackrel{\text{def}}{=} x^\top a_i^{(j)} (a_i^{(j)})^\top x$ ,  $R_j \stackrel{\text{def}}{=} \frac{1}{k} \sum_{i \in [k]} R_{i,j}$ , and let  $z$  be median value of the  $R_j$  then with probability  $1 - p$  we have that  $|z - x^\top \Sigma x| \leq \epsilon \lambda_1$ .*

*Proof.*

$$\begin{aligned} \text{Var}_{a \sim \mathcal{D}}(x^\top a a^\top x) &= \mathbb{E}_{a \sim \mathcal{D}}(x^\top a a^\top x)^2 - (\mathbb{E}_{a \sim \mathcal{D}} x^\top a a^\top x)^2 \\ &\leq \mathbb{E}_{a \sim \mathcal{D}} \|a\|_2^2 x^\top a a^\top x - (x^\top \Sigma x)^2 \\ &\leq \left\| \mathbb{E}_{a \sim \mathcal{D}} \|a\|_2^2 a a^\top \right\|_2 = v(\mathcal{D}) \lambda_1^2 \end{aligned}$$

Consequently,  $\text{Var}(R_{i,j}) \leq v(\mathcal{D})\lambda_1^2$ , and since each of the  $a_i^{(j)}$  were drawn independently this implies that we have that  $\text{Var}(R_j) \leq v(\mathcal{D})\lambda_1^2/k$ . Therefore, by Chebyshev's inequality

$$\mathbb{P} \left[ |R_j - \mathbb{E}[R_j]| \geq 2\sqrt{\frac{v(\mathcal{D})\lambda_1^2}{k}} \right] \leq \frac{1}{4}.$$

Since  $\mathbb{E}[R_j] = x^\top \Sigma x$  and since we defined  $k$  appropriately this implies that

$$\mathbb{P} \left[ |R_j - x^\top \Sigma x| \geq \epsilon \lambda_1 \right] \leq \frac{1}{4}. \quad (12)$$

The median  $z$  satisfies  $|z - x^\top \Sigma x| \leq \epsilon$  as more than half of the  $R_j$  satisfy  $|R_j - x^\top \Sigma x| \leq \epsilon$ . This happens with probability  $1 - p$  by Chernoff bound, our choice of  $m$  and (12).  $\square$

## 5.2 Solving the Linear system

Here we show how to solve linear systems in  $\mathbf{B} = \lambda \mathbf{I} - \Sigma$  in the streaming setting. We follow the general strategy of the offline algorithms given in Section 4, replacing traditional SVRG with the streaming SVRG algorithm of [FGKS14]. Again we minimize  $f(x) = \frac{1}{2}x^\top \mathbf{B}x - b^\top x$ . Similarly to in the offline case, we define for all  $a \in \text{supp}(\mathcal{D})$ ,

$$\psi_a(x) \stackrel{\text{def}}{=} \frac{1}{2}x^\top (\lambda \mathbf{I} - aa^\top)x - b^\top x. \quad (13)$$

This definition insures that

$$f(x) = \mathbb{E}_{a \sim \mathcal{D}} \psi_a(x).$$

The performance of streaming SVRG [FGKS14] will be governed by three regularity parameters. As in the offline case, we use the fact that  $f(\cdot)$  is  $\mu$ -strongly convexity for  $\mu = \lambda - \lambda_1$ . We also again require a smoothness parameter  $\bar{S}$  such that:

$$\forall x \in \mathbb{R}^d : \mathbb{E}_{a \sim \mathcal{D}} \|\nabla \psi_a(x) - \nabla \psi_a(x^{\text{opt}})\|_2^2 \leq 2\bar{S} [f(x) - f(x^{\text{opt}})]. \quad (14)$$

Lastly, we need an upper bound the variance,  $\sigma^2$  such that:

$$\mathbb{E}_{a \sim \mathcal{D}} \frac{1}{2} \|\nabla \psi_a(x^{\text{opt}})\|_{(\nabla^2 f(x^{\text{opt}}))^{-1}}^2 \leq \sigma^2 \quad (15)$$

We bound the second two parameters as follows.

**Lemma 19** (Streaming Smoothness). *The smoothness parameter*

$$\bar{S} \stackrel{\text{def}}{=} \lambda + \frac{v(\mathcal{D})\lambda_1^2}{\lambda - \lambda_1}$$

*satisfies (14).*

*Proof.* Our proof is similar to the one given in Lemma 10.

$$\begin{aligned}
\mathbb{E}_{a \sim \mathcal{D}} \|\nabla \psi_a(x) - \nabla \psi_a(x^{\text{opt}})\|_2^2 &= \mathbb{E}_{a \sim \mathcal{D}} \left\| (\lambda \mathbf{I} - aa^\top)(x - x^{\text{opt}}) \right\|_2^2 \\
&= \lambda^2 \|x - x^{\text{opt}}\|_2^2 - 2\lambda \mathbb{E}_{a \sim \mathcal{D}} \|x - x^{\text{opt}}\|_{aa^\top}^2 + \mathbb{E}_{a \sim \mathcal{D}} \|aa^\top(x - x^{\text{opt}})\|_2^2 \\
&\leq \lambda^2 \|x - x^{\text{opt}}\|_2^2 - 2\lambda \|x - x^{\text{opt}}\|_\Sigma^2 + \left\| \mathbb{E}_{a \sim \mathcal{D}} \|a\|_2^2 aa^\top \right\|_2 \cdot \|x - x^{\text{opt}}\|_2^2 \\
&\leq \lambda \|x - x^{\text{opt}}\|_{\mathbf{B}}^2 + v(\mathcal{D})\lambda_1^2 \|x - x^{\text{opt}}\|_2^2.
\end{aligned}$$

Since  $f$  is  $\lambda - \lambda_1$ -strongly convex,  $\|x - x^{\text{opt}}\|_2^2 \leq \frac{2}{\lambda - \lambda_1} [f(x) - f(x^{\text{opt}})]$ . Additionally,  $2[f(x) - f(x^{\text{opt}})] = \|x - x^{\text{opt}}\|_{\mathbf{B}}^2$ . The result follows.  $\square$

**Lemma 20** (Streaming Variance). *For  $\sigma^2 \stackrel{\text{def}}{=} \mathbb{E}_{a \sim \mathcal{D}} \frac{1}{2} \|\nabla \psi_a(x^{\text{opt}})\|_{(\nabla^2 f(x^{\text{opt}}))^{-1}}^2$  we have*

$$\sigma^2 \leq \left( \frac{v(\mathcal{D})\lambda_1^2}{\lambda - \lambda_1} \right) \|x^{\text{opt}}\|_2^2.$$

*Proof.* We have

$$\begin{aligned}
\sigma^2 &= \mathbb{E}_{a \sim \mathcal{D}} \frac{1}{2} \left\| (\lambda \mathbf{I} - aa^\top) x^{\text{opt}} - b \right\|_{\mathbf{B}^{-1}}^2 \\
&= \mathbb{E}_{a \sim \mathcal{D}} \frac{1}{2} \left\| (\lambda \mathbf{I} - aa^\top) x^{\text{opt}} - \mathbf{B} x^{\text{opt}} \right\|_{\mathbf{B}^{-1}}^2 \\
&= \mathbb{E}_{a \sim \mathcal{D}} \frac{1}{2} \left\| (\Sigma - aa^\top) x^{\text{opt}} \right\|_{\mathbf{B}^{-1}}^2.
\end{aligned}$$

Applying  $\mathbb{E} \|a - \mathbb{E}a\|_2^2 = \mathbb{E} \|a\|_2^2 - \|\mathbb{E}a\|_2^2$  gives:

$$\sigma^2 = \mathbb{E}_{a \sim \mathcal{D}} \frac{1}{2} \|x^{\text{opt}}\|_{aa^\top \mathbf{B}^{-1} aa^\top}^2 - \frac{1}{2} \|x^{\text{opt}}\|_{\Sigma \mathbf{B}^{-1} \Sigma}^2 \leq \mathbb{E}_{a \sim \mathcal{D}} \frac{1}{2} \|x^{\text{opt}}\|_{aa^\top \mathbf{B}^{-1} aa^\top}^2.$$

Furthermore, since  $\mathbf{B}^{-1} \preceq \frac{1}{\lambda - \lambda_1} \mathbf{I}$  we have

$$\mathbb{E}_{a \sim \mathcal{D}} aa^\top \mathbf{B}^{-1} aa^\top \preceq \frac{1}{\lambda - \lambda_1} \mathbb{E}_{a \sim \mathcal{D}} (aa^\top)^2 \preceq \left( \frac{\|\mathbb{E}_{a \sim \mathcal{D}} (aa^\top)^2\|_2}{\lambda - \lambda_1} \right) \mathbf{I} = \left( \frac{v(\mathcal{D})\lambda_1^2}{\lambda - \lambda_1} \right) \mathbf{I}.$$

yielding the result.  $\square$

With these two lemmas in place, we can apply the streaming SVRG algorithm of [FGKS14] to solving systems in  $\mathbf{B}$ . We encapsulate the basic iterative step of Algorithm 1 of [FGKS14] in the following definition:

**Definition 21** (Streaming SVRG Step). *Given  $x_0 \in \mathbb{R}^d$  and  $\eta, k, m > 0$  we define a streaming SVRG step,  $x = \text{ssvrg\_iter}(x_0, \eta, k, m)$  as follows. First we take  $k$  samples  $a_1, \dots, a_k$  from  $\mathcal{D}$  and set  $g = \frac{1}{k} \sum_{i \in [k]} \psi_{a_i}$  where  $\psi_{a_i}$  is as defined in (13). Then for  $\tilde{m}$  chosen uniformly at random from  $\{1, \dots, m\}$  we draw  $\tilde{m}$  additional samples  $\tilde{a}_1, \dots, \tilde{a}_{\tilde{m}}$  from  $\mathcal{D}$ . For  $t = 0, \dots, \tilde{m} - 1$  we let*

$$x_{t+1} := x_t - \frac{\eta}{L} (\nabla \psi_{\tilde{a}_t}(x_t) - \nabla \psi_{\tilde{a}_t}(x_0) + \nabla g(x_0))$$

*Finally, we return  $x_{\tilde{m}}$ .*

The accuracy of the above iterative step is proven in Theorem 4.1 of [FGKS14], which we include, using our notation below:

**Theorem 22** (Theorem 4.1 of [FGKS14]). *Letting  $f(x) = \mathbb{E}_{a \sim \mathcal{D}} \psi_a(x)$  and  $\mu, \bar{S}, \sigma^2$  be the strong convexity, smoothness, and variance bounds for  $f(x)$ , for any distribution over  $x_0$  we have that  $x := \text{ssvrg\_iter}(w_0, \eta, k, m)$  has  $\mathbb{E}[f(x) - f(x^{\text{opt}})]$  upper bounded by*

$$\frac{1}{1 - 4\eta} \left[ \left( \frac{\bar{S}}{\mu m \eta} + 4\eta \right) [f(x_0) - f(x^{\text{opt}})] + \frac{1 + 2\eta}{k} \left( \sqrt{\frac{\bar{S}}{\mu} \cdot [f(x_0) - f(x^{\text{opt}})] + \sigma} \right)^2 \right].$$

Note that Theorem 4.1 in [FGKS14] has an additional parameter of  $\alpha$ , which bounds the Hessian of  $f(\cdot)$  at the optimum in comparison to the hessian at any other point. In our setting this parameter is 1 as  $\nabla^2 f(y) = \nabla^2 f(z)$  for all  $y$  and  $z$ . With Theorem 22 we immediately have the following:

**Corollary 23** (Streaming SVRG-Based Solver). *Let  $\mu = \lambda - \lambda_1$ ,  $\bar{S} = \lambda + \frac{v(\mathcal{D})\lambda_1^2}{\lambda - \lambda_1}$ , and  $\sigma^2 = \frac{v(\mathcal{D})\lambda_1^2}{\lambda - \lambda_1} \|x^{\text{opt}}\|_2^2$ . Let  $c_2, c_3 \in (0, 1)$  be any constants,  $\eta = \frac{c_2}{8}$ ,  $m = \left\lceil \frac{\bar{S}}{\mu c_2^2} \right\rceil$ , and  $k = \max \left\{ \left\lceil \frac{\bar{S}}{\mu c_2} \right\rceil, \left\lceil \frac{v(\mathcal{D})\lambda_1^2}{(\lambda - \lambda_1)^2 c_3} \right\rceil \right\}$ . If used to solve  $\mathbf{B}x = b$  for a unit vector  $b$ , the iterative procedure described in Definition 21 returns  $x = \text{ssvrg\_iter}(x_0, \eta, k, m)$  satisfying:*

$$\mathbb{E} \|x - x^{\text{opt}}\|_{\mathbf{B}}^2 \leq 22c_2 \cdot \|x_0 - x^{\text{opt}}\|_{\mathbf{B}}^2 + 10c_3\lambda_1(\mathbf{B}^{-1}).$$

Further, the procedure requires  $O\left(\frac{v(\mathcal{D})}{\text{gap}^2} \left[\frac{1}{c_2^2} + \frac{1}{c_3}\right]\right)$  samples from  $\mathcal{D}$ .

*Proof.* Using the inequality  $(x + y)^2 \leq 2x^2 + 2y^2$  we have that

$$\left( \sqrt{\frac{\bar{S}}{\mu} \cdot \mathbb{E}[f(x_0) - f(x^{\text{opt}})] + \sigma} \right)^2 \leq \frac{2\bar{S}}{\mu} \cdot \mathbb{E}[f(x_0) - f(x^{\text{opt}})] + 2\sigma^2$$

Additionally, since  $b$  is a unit vector, we know that  $\|x^{\text{opt}}\|_2^2 = \|\mathbf{B}^{-1}b\|_2^2 \leq \frac{1}{(\lambda - \lambda_1)^2}$ . Using the fact shown in equation (10) that for any  $x$ ,  $\|x - x^{\text{opt}}\|_{\mathbf{B}}^2 = 2[f(x) - f(x^{\text{opt}})]$  we have by Theorem 22:

$$\begin{aligned} \mathbb{E} \|x - x^{\text{opt}}\|_{\mathbf{B}}^2 &\leq \frac{1}{1 - c_2/2} \left[ 8c_2 + \frac{c_2}{2} + \frac{4 + c_2}{2} \cdot c_2 \right] \cdot \|x_0 - x^{\text{opt}}\|_{\mathbf{B}}^2 \\ &\quad + \frac{4}{1 - c_2/2} \cdot \frac{4 + c_2}{4k} \cdot \frac{v(\mathcal{D})\lambda_1^2}{(\lambda - \lambda_1)^3} \\ &\leq 22c \cdot \|x_0 - x^{\text{opt}}\|_{\mathbf{B}}^2 + \frac{10c_3}{\lambda - \lambda_1} \\ &= 22c \cdot \|x_0 - x^{\text{opt}}\|_{\mathbf{B}}^2 + 10c_3\lambda_1(\mathbf{B}^{-1}). \end{aligned}$$

The number of samples required to make the streaming SVRG step is simply bounded by  $m + k$ .  $m = \frac{\bar{S}}{\mu c_2^2} = O\left(\frac{\lambda}{c_2^2(\lambda - \lambda_1)} + \frac{v(\mathcal{D})\lambda_1^2}{c_2^2(\lambda - \lambda_1)^2}\right) = O\left(\frac{1}{c_2^2 \text{gap}} + \frac{v(\mathcal{D})}{c_2^2 \text{gap}^2}\right)$ .  $\text{gap} < 1$  and  $v(\mathcal{D}) \geq 1$ . So we can simplify:  $m = O\left(\frac{v}{c_2^2 \text{gap}^2}\right)$ . We can ignore the  $\left\lceil \frac{\bar{S}}{\mu c_2} \right\rceil$  in  $k$  since this was already included in our

bound of  $m$  and so just bound  $\frac{v(\mathcal{D})\lambda_1^2}{c_3(\lambda-\lambda_1)^2} = O\left(\frac{v(\mathcal{D})}{\text{gap}^2 c_3}\right)$ . So overall the number of samples we need to take is:

$$O\left(\frac{v(\mathcal{D})}{\text{gap}^2} \left[\frac{1}{c_2^2} + \frac{1}{c_3}\right]\right).$$

□

In the offline case, when solving linear systems in the shifted-and-inverted power method, we can insure that  $\|x_0 - x^{\text{opt}}\|_{\mathbf{B}}^2$  is small by Corollary 6. In the online case, we do not have the same guarantee, since starting with a good initial value requires accurately estimating our Rayleigh quotient, which is too expensive. However, we can still show the following corollary:

**Corollary 24** (Streaming SVRG-Based Solver With No Initial Error). *There is a streaming algorithm that iteratively applies the solver of Corollary 23 to  $\mathbf{B}x = b$  for a unit vector  $b$ , which returns  $x$  satisfying:*

$$\mathbb{E} \|x - x^{\text{opt}}\|_{\mathbf{B}}^2 \leq 10c_3\lambda_1(\mathbf{B}^{-1}).$$

using  $O\left(\frac{v(\mathcal{D})}{\text{gap}^2 c_3}\right)$  samples from  $\mathcal{D}$ .

*Proof.* Let  $x_0 = 0$ . Then  $\|x_0 - x^{\text{opt}}\|_{\mathbf{B}}^2 = \|\mathbf{B}^{-1}x\|_{\mathbf{B}}^2 \leq \lambda_1(\mathbf{B}^{-1})$  since  $x$  is a unit vector. If we apply Corollary 23 with  $c_2 = 1/44$  and  $c'_3 = \frac{1}{20}$ , then we will obtain  $x_1$  with  $\mathbb{E} \|x_1 - x^{\text{opt}}\|_{\mathbf{B}}^2 \leq \frac{1}{2}\lambda_1(\mathbf{B}^{-1})$ . If we then double  $c'_3$  and apply the solver again we obtain  $x_2$  with  $\mathbb{E} \|x_1 - x^{\text{opt}}\|_{\mathbf{B}}^2 \leq \frac{1}{4}\lambda_1(\mathbf{B}^{-1})$ . Iterating in this way, after  $\log(1/c_3)$  iterations we will have the desired guarantee:  $\mathbb{E} \|x - x^{\text{opt}}\|_{\mathbf{B}}^2 \leq 10c_3\lambda_1(\mathbf{B}^{-1})$ . Our total sample cost in each iteration is, by Corollary 23,  $O\left(\frac{v(\mathcal{D})}{\text{gap}^2} \left[\frac{1}{44^2} + \frac{1}{c'_3}\right]\right)$ . Since we double  $c'_3$  each time, the cost corresponding to the  $\frac{1}{c'_3}$  terms is dominated by the last iteration when we have  $c'_3 = O(c_3)$ . So our overall sample cost is just:

$$O\left(\frac{v(\mathcal{D})}{\text{gap}^2} \left[\frac{1}{c_3} + \log(1/c_3)\right]\right) = O\left(\frac{v(\mathcal{D})}{\text{gap}^2 c_3}\right).$$

□

### 5.3 Online Shifted-and-Inverted Power Method

We now apply the results in Section 5.1 and Section 5.2 to the power method framework of Section 3 to give our main online result. We give an algorithm that quickly refines a coarse approximation to  $v_1$  into a finer approximation.

**Theorem 25** (Online Shifted-and-Inverted Power Method – Warm Start). *Let  $\mathbf{B} = \lambda\mathbf{I} - \mathbf{A}^\top \mathbf{A}$  for  $(1 + \frac{\text{gap}}{150})\lambda_1 \leq \lambda \leq (1 + \frac{\text{gap}}{100})\lambda_1$  and let  $x_0$  be some vector with  $G(x_0) \leq \frac{1}{\sqrt{10}}$ . Running the inverted power method on  $\mathbf{B}$  initialized with  $x_0$ , using the streaming SVRG solver described in Definition 21 to approximately apply  $\mathbf{B}^{-1}$  at each step, returns  $x$  such that, for any parameter  $\delta > \frac{1}{d^{10}}$  with probability  $1 - \delta$ ,  $x^\top \Sigma x \geq (1 - \epsilon)\lambda_1$  using total sample count:*

$$O\left(\frac{v(\mathcal{D})}{\text{gap}} \cdot \left(\frac{\log 1/\delta + \log \log 1/\epsilon}{\text{gap}} + \frac{1}{\delta^2 \epsilon}\right)\right)$$

The amortized processing time per sample is simple  $O(d)$ .

We note that by instantiating Theorem 25, with  $\epsilon' = \epsilon \cdot \text{gap}$ , and applying Lemma 1 we can find  $x$  such that  $|v_1^\top x| \geq 1 - \epsilon$  in time  $O\left(\frac{v(\mathcal{D})}{\text{gap}^2 \cdot \epsilon \cdot \delta^2}\right)$ .

*Proof.* By Lemma 3 it suffices to have  $G^2(x) = O(\frac{\epsilon}{\text{gap}})$  so  $G(x) = O(\sqrt{\epsilon/\text{gap}})$ . In order to have this with probability  $1 - \delta$  it suffices to have  $\mathbb{E}[G(x)] = O(\delta\sqrt{\epsilon/\text{gap}})$ . Since we start with  $G(x_0) \leq \frac{1}{\sqrt{10}}$ , we can achieve this using  $\log(\text{gap}/(\delta\epsilon))$  iterations of the approximate power method of Theorem 5. In iteration  $i$  we choose the error parameter for Theorem 5 to be  $c_1(i) = \frac{1}{\sqrt{10}} \cdot \left(\frac{1}{5}\right)^i$ . In this way, we have:

$$\mathbb{E}[G(x_i)] \leq \frac{3}{25}G(x_{i-1}) + \frac{4}{1000} \frac{1}{\sqrt{10}} \cdot \left(\frac{1}{5}\right)^i$$

and by induction  $\mathbb{E}[G(x_i)] \leq \frac{1}{5^i} \frac{1}{\sqrt{10}}$ . We halt when  $\left(\frac{1}{5}\right)^i = O(\delta\sqrt{\epsilon/\text{gap}})$  and hence  $c_1(i) = O(\delta\sqrt{\epsilon/\text{gap}})$ .

In order to apply Theorem 5 we need a subroutine  $\widehat{\text{quot}}(x)$  that lets us approximate  $\text{quot}(x)$  to within an additive error  $\frac{1}{30}(\lambda - \lambda_1) = O(\text{gap}\lambda_1)$ . Theorem 18 gives us such a routine, requiring  $O\left(\frac{v(\mathcal{D}) \log(1/\delta \log(\text{gap}/\delta\epsilon))}{\text{gap}^2}\right)$  samples to succeed with probability  $1 - O\left(\frac{\delta}{\log(\text{gap}/\delta\epsilon)}\right)$ . Union bounding, the estimation succeeds in all rounds of the power method with probability at least  $1 - O(\delta)$ .

Our cost for each linear system solve is given by Corollary 24 with  $c_3 = \Theta(c_1(i)^2)$ , is:

$$O\left(\frac{v(\mathcal{D})}{\text{gap}^2 c_1(i)^2}\right).$$

Now,  $c_1(i)$  multiplies by a constant factor with each iteration. The cost over all  $O(\log(\text{gap}/\delta\epsilon))$  iterations is just a truncated geometric series and is proportional to cost in the last iteration, when  $c_3 = \Theta\left(\frac{\delta^2 \epsilon}{\text{gap}}\right)$ . So the total cost for solving the linear systems is

$$O\left(\frac{v(\mathcal{D})}{\text{gap} \delta^2 \epsilon}\right).$$

Adding this to the Rayleigh quotient estimation cost give us total sample count:

$$O\left(\frac{v(\mathcal{D})}{\text{gap}^2} \cdot \log(1/\delta \log(1/\epsilon)) + \frac{v(\mathcal{D})}{\text{gap} \delta^2 \epsilon}\right) = O\left(\frac{v(\mathcal{D})}{\text{gap}} \cdot \left(\frac{\log 1/\delta + \log \log 1/\epsilon}{\text{gap}} + \frac{1}{\delta^2 \epsilon}\right)\right).$$

□

## 6 Parameter Estimation for Offline Eigenvector Computation

In Section 4, in order to invoke Theorems 5 and 8 we assumed knowledge of some  $\lambda$  with  $(1 + c_1 \cdot \text{gap})\lambda_1 \leq \lambda \leq (1 + c_2 \cdot \text{gap})\lambda_1$  for some small constant  $c_1$  and  $c_2$  ( $\hat{\lambda}_1$  in Theorem 5 could also be obtained in this form). Here we show how to estimate this parameter using Algorithm 1, completing the proof of our offline eigenvector estimation algorithm (Theorem 16).

In this section, for simplicity we assume that we have oracle access to compute  $\mathbf{B}_\lambda^{-1}w$  for any given  $w$ , and any  $\lambda > \lambda_1$ , but the results here can be extended to the case where we can compute  $\mathbf{B}_\lambda^{-1}w$  only approximately. We will use a result of [MM15] that gives gap free bounds for computing eigenvalues using the power method. The following is a specialization of Theorem 1 from [MM15]:

---

**Algorithm 1** Estimating the eigenvalue and the eigengap

---

**Input:**  $\mathbf{A} \in \mathbb{R}^{n \times d}$ ,  $\alpha$

- 1:  $[w_1, w_2] \leftarrow \mathcal{N}(0, I^{d \times d})$
- 2:  $t \leftarrow O(\alpha \log d)$
- 3:  $[\tilde{\lambda}_1^{(0)}, \tilde{\lambda}_2^{(0)}] \leftarrow \text{eigEstimate}\left((\mathbf{A}^T \mathbf{A})^t w\right)$
- 4:  $\bar{\lambda}^{(0)} \leftarrow (1 + \frac{1}{2})\tilde{\lambda}_1^{(0)}$
- 5:  $i \leftarrow 0$
- 6: **while**  $\bar{\lambda}^{(i)} - \tilde{\lambda}_1^{(i)} < \frac{1}{10}(\bar{\lambda}^{(i)} - \tilde{\lambda}_2^{(i)})$  **do**
- 7:    $i \leftarrow i + 1$
- 8:    $[w_1, w_2] \leftarrow \mathcal{N}(0, I^{d \times d})$
- 9:    $[\hat{\lambda}_1^{(i)}, \hat{\lambda}_2^{(i)}] \leftarrow \text{eigEstimate}\left(\left(\bar{\lambda}^{(i-1)} \mathbf{I} - \mathbf{A}^T \mathbf{A}\right)^{-t} w\right)$
- 10:    $[\tilde{\lambda}_1^{(i)}, \tilde{\lambda}_2^{(i)}] \leftarrow \left[\bar{\lambda}^{(i-1)} - \frac{1}{\hat{\lambda}_1^{(i)}}, \bar{\lambda}^{(i-1)} - \frac{1}{\hat{\lambda}_2^{(i)}}\right]$
- 11:    $\bar{\lambda}^{(i)} \leftarrow \frac{1}{2}(\tilde{\lambda}_1^{(i)} + \bar{\lambda}^{(i-1)})$
- 12: **end while**

**Output:**  $\lambda$

---

**Theorem 26.** For any  $\epsilon > 0$ , any matrix  $\mathbf{M} \in \mathbb{R}^{d \times d}$  with eigenvalues  $\lambda_1, \dots, \lambda_d$ , and  $k \leq d$ , let  $\mathbf{W} \in \mathbb{R}^{d \times k}$  be a matrix with entries drawn independently from  $\mathcal{N}(0, 1)$ . Let  $\text{eigEstimate}(\mathbf{Y})$  be a function returning for each  $i$ ,  $\tilde{\lambda}_i = \tilde{v}_i^T \mathbf{M} \tilde{v}_i$  where  $\tilde{v}_i$  is the  $i^{\text{th}}$  largest eigenvector of  $\mathbf{Y}$ . Then setting  $[\tilde{\lambda}_1, \dots, \tilde{\lambda}_k] = \text{eigEstimate}(\mathbf{M}^t \mathbf{W})$ , for some fixed constant  $c$  and  $t = c\alpha \log d$  for any  $\alpha > 1$ , we have for all  $i$ :

$$|\tilde{\lambda}_i - \lambda_i| \leq \frac{1}{\alpha} \lambda_{k+1}$$

Throughout the proof, we assume  $\alpha$  is picked to be some large constant so that  $\alpha > 100$ , then Theorem 26 implies:

**Lemma 27.** The iterates of Algorithm 1 satisfy:

$$0 \leq \lambda_1 - \tilde{\lambda}_1^{(0)} \leq \frac{1}{\alpha} \lambda_1 \quad \text{and} \quad \frac{1}{2} \left(1 - \frac{3}{\alpha}\right) \lambda_1 \leq \bar{\lambda}^{(0)} - \lambda_1 \leq \frac{1}{2} \lambda_1, \quad \text{and},$$

$$0 \leq \lambda_1 - \tilde{\lambda}_1^{(i)} \leq \frac{1}{\alpha - 1} (\bar{\lambda}^{(i-1)} - \lambda_1) \quad \text{and} \quad \frac{1}{2} \left(1 - \frac{1}{\alpha - 1}\right) (\bar{\lambda}^{(i-1)} - \lambda_1) \leq \bar{\lambda}^{(i)} - \lambda_1 \leq \frac{1}{2} (\bar{\lambda}^{(i-1)} - \lambda_1).$$

*Proof.* The proof can be decomposed into two parts:

**Part I (Lines 3-4):** Theorem 26 tells us that  $\tilde{\lambda}_1^{(0)} \geq (1 - \frac{1}{\alpha}) \lambda_1$ . This means that we have

$$0 \leq \lambda_1 - \tilde{\lambda}_1^{(0)} \leq \frac{1}{\alpha} \lambda_1 \quad \text{and} \quad \frac{1}{2} \left(1 - \frac{3}{\alpha}\right) \lambda_1 \leq \bar{\lambda}^{(0)} - \lambda_1 \leq \frac{1}{2} \lambda_1.$$



**Part II (Lines 5-6):** Consider now iteration  $i$ . We now apply Theorem 26 to the matrix  $(\bar{\lambda}^{(i-1)} \mathbf{I} - \mathbf{A}^T \mathbf{A})^{-1}$ . The top eigenvalue of this matrix is  $(\bar{\lambda}^{(i-1)} - \lambda_1)^{-1}$ . This means that we have  $(1 - \frac{1}{\alpha}) (\bar{\lambda}^{(i-1)} - \lambda_1)^{-1} \leq \hat{\lambda}_1^{(i)} \leq (\bar{\lambda}^{(i-1)} - \lambda_1)^{-1}$ , and hence we have,

$$0 \leq \lambda_1 - \tilde{\lambda}_1^{(i)} \leq \frac{1}{\alpha - 1} (\bar{\lambda}^{(i-1)} - \lambda_1) \quad \text{and} \quad \frac{1}{2} \left(1 - \frac{1}{\alpha - 1}\right) (\bar{\lambda}^{(i-1)} - \lambda_1) \leq \bar{\lambda}^{(i)} - \lambda_1 \leq \frac{1}{2} (\bar{\lambda}^{(i-1)} - \lambda_1).$$

This proves the lemma.  $\square$

**Lemma 28.** Recall we denote  $\lambda_2 \stackrel{\text{def}}{=} \lambda_2(\mathbf{A}^T \mathbf{A})$  and  $\text{gap} \stackrel{\text{def}}{=} \frac{\lambda_1 - \lambda_2}{\lambda_1}$ . The iterates of Algorithm 1 satisfy  $|\lambda_2 - \tilde{\lambda}_2^{(i)}| \leq \frac{1}{\alpha - 1} (\bar{\lambda}^{(i-1)} - \lambda_2)$ , and  $\bar{\lambda}^{(i)} - \tilde{\lambda}_2^{(i)} \geq \frac{\text{gap} \lambda_1}{4}$ .

*Proof.* Since  $(\bar{\lambda}^{(i-1)} - \lambda_2)^{-1}$  is the second eigenvalue of the matrix  $(\bar{\lambda}^{(i-1)} \mathbf{I} - \mathbf{A}^T \mathbf{A})^{-1}$ , Theorem 26 tells us that

$$\left(1 - \frac{1}{\alpha}\right) (\bar{\lambda}^{(i-1)} - \lambda_2)^{-1} \leq \hat{\lambda}_2^{(i)} \leq \left(1 + \frac{1}{\alpha}\right) (\bar{\lambda}^{(i-1)} - \lambda_2)^{-1}.$$

This immediately yields the first claim. For the second claim, we notice that

$$\begin{aligned} \bar{\lambda}^{(i)} - \tilde{\lambda}_2^{(i)} &= \bar{\lambda}^{(i)} - \lambda_2 + \lambda_2 - \tilde{\lambda}_2^{(i)} \\ &\stackrel{(\zeta_1)}{\geq} \bar{\lambda}^{(i)} - \lambda_2 - \frac{1}{\alpha - 1} (\bar{\lambda}^{(i-1)} - \lambda_2) \\ &= \bar{\lambda}^{(i)} - \lambda_1 - \frac{1}{\alpha - 1} (\bar{\lambda}^{(i-1)} - \lambda_1) + \left(1 - \frac{1}{\alpha - 1}\right) (\lambda_1 - \lambda_2) \\ &\stackrel{(\zeta_2)}{\geq} \frac{1}{2} \left(1 - \frac{3}{\alpha - 1}\right) (\bar{\lambda}^{(i-1)} - \lambda_1) + \left(1 - \frac{1}{\alpha - 1}\right) (\lambda_1 - \lambda_2) \geq \frac{\text{gap} \lambda_1}{4}, \end{aligned}$$

where  $(\zeta_1)$  follows from the first claim of this lemma, and  $(\zeta_2)$  follows from Lemma 27.  $\square$

We now state and prove the final result:

**Theorem 29.** Suppose  $\alpha > 100$ , and after  $T$  iterations, Algorithm 1 exits (i.e.  $T$  is the final value of iterates  $i$  in Algorithm), then we have  $T \leq \left\lceil \log \frac{10}{\text{gap}} \right\rceil + 1$ , and we will have:

$$\left(1 + \frac{\text{gap}}{120}\right) \lambda_1 \leq \bar{\lambda}^{(T)} \leq \left(1 + \frac{\text{gap}}{8}\right) \lambda_1$$

*Proof.* Let  $\bar{i} = \left\lceil \log \frac{10}{\text{gap}} \right\rceil$ , suppose the algorithm has not exited yet after  $\bar{i}$  iterations, then since  $\bar{\lambda}^{(i)} - \lambda_1$  decays geometrically, we have  $\bar{\lambda}^{(\bar{i})} - \lambda_1 \leq \frac{\text{gap} \lambda_1}{10}$ . Therefore, Lemmas 27 and 28 imply that  $\bar{\lambda}^{(\bar{i}+1)} - \tilde{\lambda}_1^{(\bar{i}+1)} \leq \left(\frac{1}{2} + \frac{1}{\alpha - 1}\right) (\bar{\lambda}^{(\bar{i})} - \lambda_1) \leq \frac{\text{gap} \lambda_1}{15}$ , and

$$\begin{aligned} \bar{\lambda}^{(\bar{i}+1)} - \tilde{\lambda}_2^{(\bar{i}+1)} &\geq \bar{\lambda}^{(\bar{i}+1)} - \lambda_2 - \left|\lambda_2 - \tilde{\lambda}_2^{(\bar{i}+1)}\right| \geq \lambda_1 - \lambda_2 - \frac{1}{\alpha - 1} (\bar{\lambda}^{(\bar{i})} - \lambda_2) \\ &= \text{gap} \lambda_1 - \frac{1}{\alpha - 1} (\bar{\lambda}^{(\bar{i})} - \lambda_1 + \lambda_1 - \lambda_2) \geq \frac{3}{4} \text{gap} \lambda_1 \end{aligned}$$

This means that the exit condition on Line 6 must be triggered in  $\bar{i} + 1$  iteration, proving the first part of the lemma.

For upper bound, by Lemmas 27, 28 and exit condition we know:

$$\begin{aligned}
\bar{\lambda}^{(T)} - \lambda_1 &\leq \bar{\lambda}^{(T)} - \tilde{\lambda}_1^{(T)} \leq \frac{1}{10}(\bar{\lambda}^{(T)} - \tilde{\lambda}_2^{(T)}) \leq \frac{1}{10} \left( \bar{\lambda}^{(T)} - \lambda_2 + |\lambda_2 - \tilde{\lambda}_2^{(T)}| \right) \\
&\leq \frac{1}{10} \left( \bar{\lambda}^{(T)} - \lambda_2 + \frac{1}{\alpha - 1}(\bar{\lambda}^{(T-1)} - \lambda_2) \right) \\
&= \frac{1}{10} \left( \frac{\alpha}{\alpha - 1} \text{gap} \lambda_1 + (\bar{\lambda}^{(T)} - \lambda_1) + \frac{1}{\alpha - 1} (\bar{\lambda}^{(T-1)} - \lambda_1) \right) \\
&\leq \frac{1}{10} \left( \frac{\alpha}{\alpha - 1} \text{gap} \lambda_1 + \frac{\alpha}{\alpha - 2} (\bar{\lambda}^{(T)} - \lambda_1) \right)
\end{aligned}$$

Since  $\alpha > 100$ , this directly implies  $\bar{\lambda}^{(T)} - \lambda_1 \leq \frac{\text{gap}}{8} \lambda_1$ .

For lower bound, since as long as the Algorithm 1 does not exists, by Lemmas 28, we have  $\bar{\lambda}^{(T-1)} - \tilde{\lambda}_1^{(T-1)} \geq \frac{1}{10} (\bar{\lambda}^{(T-1)} - \tilde{\lambda}_2^{(T-1)}) \geq \frac{\text{gap} \lambda_1}{40}$ , and thus:

$$\begin{aligned}
\bar{\lambda}^{(T-1)} - \lambda_1 &= \bar{\lambda}^{(T-1)} - \tilde{\lambda}_1^{(T-1)} - (\lambda_1 - \tilde{\lambda}_1^{(T-1)}) \geq \frac{\text{gap} \lambda_1}{40} - \frac{1}{\alpha - 1} (\bar{\lambda}^{(T-1)} - \lambda_1) \\
&\geq \frac{\text{gap} \lambda_1}{40} - \frac{2}{\alpha - 2} (\bar{\lambda}^{(T)} - \lambda_1) \geq \frac{\text{gap} \lambda_1}{50}
\end{aligned}$$

By Lemma 27, we know  $\bar{\lambda}^{(T)} - \lambda_1 \geq \frac{1}{2} (1 - \frac{1}{\alpha - 1} (\bar{\lambda}^{(T-1)} - \lambda_1)) > \frac{\text{gap}}{120} \lambda_1$   $\square$

Note that, although we proved the upper bound and lower bound in Theorem 29 with specific constants coefficient  $\frac{1}{8}$  and  $\frac{1}{120}$ , this analysis can easily be extended to any smaller constants by modifying the constant in exit condition, and choosing  $\alpha$  larger.

## 7 Lower Bounds

Here we show that our online eigenvector estimation algorithm (Theorem 25) is asymptotically optimal - as sample size grows large it achieves optimal accuracy as a function of sample size. We rely on the following lower bound for eigenvector estimation in the Gaussian spike model:

**Lemma 30** (Lower bound for Gaussian Spike Model [BJNP13]). *Suppose data is generated as*

$$a_i = \sqrt{\lambda} \iota_i v^* + Z_i \quad (16)$$

where  $\iota_i \sim \mathcal{N}(0, 1)$ , and  $Z_i \sim \mathcal{N}(0, I_d)$ . Let  $\hat{v}$  be some estimator of top eigenvector  $v^*$ . Then, there exist some universal constant  $c_0$ , so that for  $n$  sufficiently large, we have:

$$\inf_{\hat{v}} \max_{v^* \in \mathbb{S}^{d-1}} \mathbb{E} \|\hat{v} - v^*\|_2 \geq c_0 \frac{(1 + \lambda)d}{\lambda^2 n}$$

**Theorem 31.** *Consider the problem of estimating the top eigenvector  $v_1$  of  $\mathbb{E}_{a \sim \mathcal{D}} a a^\top$ , where we observe  $n$  i.i.d samples from unknown distribution  $\mathcal{D}$ . If  $\text{gap} < 0.9$ , then there exist some universal*

constant  $c$ , for any estimator  $\hat{v}$  of top eigenvector, there always exists some hard distribution  $\mathcal{D}$  so that for  $n$  sufficiently large:

$$\mathbb{E} \|\hat{v} - v_1\|_2^2 \geq c \frac{v(\mathcal{D})}{\text{gap}^2 n}$$

*Proof.* Suppose the claim of theorem is not true, then there exist some estimator  $\hat{v}$  so that

$$\mathbb{E} \|\hat{v} - v_1\|_2^2 < c' \frac{v(\mathcal{D})}{\text{gap}^2 n}$$

holds for all distribution  $\mathcal{D}$ , and for any fixed constant  $c'$  when  $n$  is sufficiently large.

Let distribution  $\mathcal{D}$  be the Gaussian Spike Model specified by Eq.(16), then by calculation, it's not hard to verify that:

$$v(\mathcal{D}) = \frac{\left\| \mathbb{E}_{a \sim \mathcal{D}} \left[ (aa^\top)^2 \right] \right\|_2}{\left\| \mathbb{E}_{a \sim \mathcal{D}} (aa^\top) \right\|_2^2} = \frac{d + 2 + 3\lambda}{1 + \lambda}$$

Since we know  $\text{gap} = \frac{\lambda}{1+\lambda} < 0.9$ , this implies  $\lambda < 9$ , which gives  $v(\mathcal{D}) < \frac{d+29}{1+\lambda} < \frac{30d}{1+\lambda}$ . Therefore, we have:

$$\mathbb{E} \|\hat{v} - v^*\|_2^2 < c' \frac{v(\mathcal{D})}{\text{gap}^2 n} < 30c' \frac{(1+\lambda)d}{\lambda^2 n}$$

holds for all  $v^* \in \mathbb{S}^{d-1}$ . Choose  $c' = \frac{c_0}{30}$  in Lemma 30 we have a contradiction.  $\square$

$\|\hat{v} - v_1\|_2^2 = 2 - 2\hat{v}^\top v_1$ , so this bound implies that- to obtain  $|\hat{v}^\top v_1| \geq 1 - \epsilon$ , we need  $\frac{v(\mathcal{D})}{\text{gap}^2 n} = O(\epsilon)$  so  $n = \Theta\left(\frac{v(\mathcal{D})}{\text{gap}^2 \epsilon}\right)$ . This exactly matches the sample complexity given by Theorem 25.

## 8 Acknowledgements

We would like to thank Dan Garber and Elad Hazan for helpful discussions on our work.

## References

- [AC09] Nir Ailon and Bernard Chazelle. The fast johnson-lindenstrauss transform and approximate nearest neighbors. *SIAM Journal on Computing*, 39(1):302–322, 2009.
- [BDF13] Akshay Balsubramani, Sanjoy Dasgupta, and Yoav Freund. The fast convergence of incremental pca. In *Advances in Neural Information Processing Systems*, pages 3174–3182, 2013.
- [BJNP13] Aharon Birnbaum, Iain M Johnstone, Boaz Nadler, and Debashis Paul. Minimax bounds for sparse pca with noisy high-dimensional data. *Annals of statistics*, 41(3):1055, 2013.
- [Bot10] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010.

- [CLM<sup>+</sup>15] Michael B Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, Richard Peng, and Aaron Sidford. Uniform sampling for matrix approximation. In *Proceedings of the 6th Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 181–190, 2015.
- [CR15] Dominik Csiba and Peter Richtárik. Primal method for erm with flexible mini-batching schemes and non-convex losses. *arXiv preprint arXiv:1506.02227*, 2015.
- [CW13] Kenneth L Clarkson and David P Woodruff. Low rank approximation and regression in input sparsity time. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 81–90. ACM, 2013.
- [FGKS14] Roy Frostig, Rong Ge, Sham M Kakade, and Aaron Sidford. Competing with the empirical risk minimizer in a single pass. *arXiv preprint arXiv:1412.6606*, 2014.
- [FGKS15] Roy Frostig, Rong Ge, Sham M Kakade, and Aaron Sidford. Un-regularizing: approximate proximal point and faster stochastic algorithms for empirical risk minimization. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, 2015.
- [GH15] Dan Garber and Elad Hazan. Fast and simple pca via convex optimization. *Computing Research Repository (CoRR)*, abs/1509.05647, 2015.
- [GHJ<sup>+</sup>16] Dan Garber, Elad Hazan, Chi Jin, Sham M. Kakade, Cameron Musco, Praneeth Netrapalli, and Aaron Sidford. Faster eigenvector computation via shift-and-invert preconditioning. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, 2016. Full version at [arXiv:1605.08754](https://arxiv.org/abs/1605.08754).
- [HP14] Moritz Hardt and Eric Price. The noisy power method: A meta algorithm with applications. In *Advances in Neural Information Processing Systems*, pages 2861–2869, 2014.
- [Jol02] Ian Jolliffe. *Principal component analysis*. Wiley Online Library, 2002.
- [JZ13] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems 26 (NIPS)*, pages 315–323, 2013.
- [Kor03] Yehuda Koren. On spectral graph drawing. In *Computing and Combinatorics*, pages 496–508. Springer, 2003.
- [LG14] François Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th international symposium on symbolic and algebraic computation*, pages 296–303. ACM, 2014.
- [LMH15] Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. *arXiv preprint arXiv:1506.02186*, 2015.
- [MCJ13] Ioannis Mitliagkas, Constantine Caramanis, and Prateek Jain. Memory limited, streaming pca. In *Advances in Neural Information Processing Systems*, pages 2886–2894, 2013.

- [MM15] Cameron Musco and Christopher Musco. Stronger approximate singular value decomposition via the block lanczos and power methods. *arXiv preprint arXiv:1504.05477*, 2015.
- [NJW02] Andrew Y Ng, Michael I Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems 15 (NIPS)*, pages 849–856, 2002.
- [PBMW99] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: bringing order to the web. 1999.
- [Saa92] Youcef Saad. *Numerical methods for large eigenvalue problems*, volume 158. SIAM, 1992.
- [Sha15a] Ohad Shamir. Convergence of stochastic gradient descent for pca. *arXiv preprint arXiv:1509.09002*, 2015.
- [Sha15b] Ohad Shamir. Fast stochastic algorithms for svd and pca: Convergence properties and convexity. *arXiv preprint arXiv:1507.08788*, 2015.
- [Sha15c] Ohad Shamir. A stochastic pca and svd algorithm with an exponential convergence rate. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pages 144–152, 2015.
- [Spi07] Daniel A Spielman. Spectral graph theory and its applications. In *null*, pages 29–38. IEEE, 2007.
- [SRO15] Christopher D Sa, Christopher Re, and Kunle Olukotun. Global convergence of stochastic gradient descent for some non-convex matrix problems. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 2332–2341, 2015.
- [SS15] Shai Shalev-Shwartz. Sdca without duality. *arXiv preprint arXiv:1502.06177*, 2015.
- [Tro15] Joel A Tropp. An introduction to matrix concentration inequalities. *arXiv preprint arXiv:1501.01571*, 2015.
- [VW04] Santosh Vempala and Grant Wang. A spectral algorithm for learning mixture models. *Journal of Computer and System Sciences*, 68(4):841–860, 2004.
- [Wil12] Virginia Vassilevska Williams. Multiplying matrices faster than coppersmith-winograd. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 887–898. ACM, 2012.

## A Appendix

**Lemma 32** (Eigenvector Estimation via Spectral Norm Matrix Approximation). *Let  $\mathbf{A}^\top \mathbf{A}$  have top eigenvector  $1$ , top eigenvector  $v_1$  and eigenvalue gap  $\text{gap}$ . Let  $\mathbf{B}^\top \mathbf{B}$  be some matrix with  $\|\mathbf{A}^\top \mathbf{A} - \mathbf{B}^\top \mathbf{B}\|_2 \leq O(\sqrt{\epsilon} \cdot \text{gap})$ . Let  $x$  be the top eigenvector of  $\mathbf{B}^\top \mathbf{B}$ . Then:*

$$|x^\top v_1| \geq 1 - \epsilon .$$

*Proof.* We can any unit vector  $y$  as  $y = c_1 v_1 + c_2 v_2$  where  $v_2$  is the component of  $x$  orthogonal to  $v_1$  and  $c_1^2 + c_2^2 = 1$ . We know that

$$\begin{aligned} v_1^\top \mathbf{B}^\top \mathbf{B} v_1 &= v_1^\top \mathbf{A}^\top \mathbf{A} v_1 - v_1^\top (\mathbf{A}^\top \mathbf{A} - \mathbf{B}^\top \mathbf{B}) v_1 \\ 1 - \sqrt{\epsilon} \text{gap} &\leq v_1^\top \mathbf{B}^\top \mathbf{B} v_1 \leq 1 + \sqrt{\epsilon} \text{gap} \end{aligned}$$

Similarly we can compute:

$$\begin{aligned} v_2^\top \mathbf{B}^\top \mathbf{B} v_2 &= v_2^\top \mathbf{A}^\top \mathbf{A} v_2 - v_2^\top (\mathbf{A}^\top \mathbf{A} - \mathbf{B}^\top \mathbf{B}) v_2 \\ 1 - \text{gap} - \sqrt{\epsilon} \text{gap} &\leq v_2^\top \mathbf{B}^\top \mathbf{B} v_2 \leq 1 - \text{gap} + \sqrt{\epsilon} \text{gap}. \end{aligned}$$

and

$$\begin{aligned} |v_1^\top \mathbf{B}^\top \mathbf{B} v_2| &= |v_1^\top \mathbf{A}^\top \mathbf{A} v_2 - v_1^\top (\mathbf{A}^\top \mathbf{A} - \mathbf{B}^\top \mathbf{B}) v_2| \\ &\leq \sqrt{\epsilon} \text{gap}. \end{aligned}$$

We have  $x^\top \mathbf{B} \mathbf{B}^\top x = c_1^2 (v_1^\top \mathbf{B}^\top \mathbf{B} v_1) + c_2^2 (v_2^\top \mathbf{B}^\top \mathbf{B} v_2) + 2c_1 c_2 \cdot v_2^\top \mathbf{B}^\top \mathbf{B} v_1$ .

We want to bound  $c_1 \geq 1 - \epsilon$  so  $c_1^2 \geq 1 - O(\epsilon)$ . Since  $x$  is the top eigenvector of  $\mathbf{B} \mathbf{B}^\top$  we have:

$$\begin{aligned} x^\top \mathbf{B} \mathbf{B}^\top x &\geq v_1^\top \mathbf{B} \mathbf{B}^\top v_1 \\ c_2^2 (v_2^\top \mathbf{B}^\top \mathbf{B} v_2) + 2c_2 v_2^\top \mathbf{B}^\top \mathbf{B} v_1 &\geq (1 - c_1^2) v_1^\top \mathbf{B} \mathbf{B}^\top v_1 \\ 2\sqrt{1 - c_1^2} \sqrt{\epsilon} \text{gap} &\geq (1 - c_1^2) (v_1^\top \mathbf{B} \mathbf{B}^\top v_1 - v_2^\top \mathbf{B} \mathbf{B}^\top v_2) \\ \frac{1}{\sqrt{1 - c_1^2}} &\geq \frac{(1 - 2\sqrt{\epsilon}) \text{gap}}{2\sqrt{\epsilon} \text{gap}} \\ \frac{1}{1 - c_1^2} &\geq \frac{1 - 5\sqrt{\epsilon}}{4\epsilon} \end{aligned}$$

This means we need have  $1 - c_1^2 \leq O(\epsilon)$  meaning  $c_1^2 \geq 1 - O(\epsilon)$  as desired.  $\square$

**Lemma 33** (Inverted Power Method progress in  $\ell_2$  and  $\mathbf{B}$  norms). *Let  $x$  be a unit vector with  $\langle x, v_1 \rangle \neq 0$  and let  $\tilde{x} = \mathbf{B}^{-1}x$ , i.e. the power method update of  $\mathbf{B}^{-1}$  on  $x$ . Then, we have both:*

$$\frac{\|\mathbf{P}_{v_1^\perp} \tilde{x}\|_{\mathbf{B}}}{\|\mathbf{P}_{v_1} \tilde{x}\|_{\mathbf{B}}} \leq \frac{\lambda_2(\mathbf{B}^{-1})}{\lambda_1(\mathbf{B}^{-1})} \cdot \frac{\|\mathbf{P}_{v_1^\perp} x\|_{\mathbf{B}}}{\|\mathbf{P}_{v_1} x\|_{\mathbf{B}}} \quad (17)$$

and

$$\frac{\|\mathbf{P}_{v_1^\perp} \tilde{x}\|_2}{\|\mathbf{P}_{v_1} \tilde{x}\|_2} \leq \frac{\lambda_2(\mathbf{B}^{-1})}{\lambda_1(\mathbf{B}^{-1})} \cdot \frac{\|\mathbf{P}_{v_1^\perp} x\|_2}{\|\mathbf{P}_{v_1} x\|_2} \quad (18)$$

*Proof.* (17) was already shown in Lemma 4. We show (18) similarly.

Writing  $x$  in the eigenbasis of  $\mathbf{B}^{-1}$ , we have  $x = \sum_i \alpha_i v_i$  and  $\tilde{x} = \sum_i \alpha_i \lambda_i (\mathbf{B}^{-1}) v_i$ . Since  $\langle x, v_1 \rangle \neq 0$ ,  $\alpha_1 \neq 0$  and we have:

$$\frac{\|\mathbf{P}_{v_1^\perp} \tilde{x}\|_2}{\|\mathbf{P}_{v_1} \tilde{x}\|_2} = \frac{\sqrt{\sum_{i \geq 2} \alpha_i^2 \lambda_i^2 (\mathbf{B}^{-1})}}{\sqrt{\alpha_1^2 \lambda_1^2 (\mathbf{B}^{-1})}} \leq \frac{\lambda_2 (\mathbf{B}^{-1})}{\lambda_1 (\mathbf{B}^{-1})} \cdot \frac{\sqrt{\sum_{i \geq 2} \alpha_i^2}}{\sqrt{\alpha_1^2}} = \frac{\lambda_2 (\mathbf{B}^{-1})}{\lambda_1 (\mathbf{B}^{-1})} \cdot \frac{\|\mathbf{P}_{v_1^\perp} x\|_2}{\|\mathbf{P}_{v_1} x\|_2}.$$

$\square$